NOTICE

AVIS

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

If pages are missing, contact the university which granted the degree.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30, and subsequent amendments.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30, et ses amendements subséquents.

Canada

AN EPISODIC-PROCESSING ACCOUNT OF IMPLICIT GRAMMAR LEARNING:

PRESERVATION OF THE FUNCTIONAL OPERATIONS PERFORMED ON

PARTICULAR ASPECTS OF ITEM STRUCTURE

by

Michael Douglas Dorken

B. Arts & Sci. (Honours), McMaster, 1990

THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR

THE DEGREE OF MASTER OF ARTS

in the Department

of

Psychology

© Michael Douglas Dorken 1992

SIMON FRASER UNIVERSITY

August, 1992

The author has granted an irrevocable non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of his/her thesis by any means and in any form or format, making this thesis available to interested persons.

The author retains ownership of the copyright in his/her thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without his/her permission.

L'auteur a accordé une licence irrévocable et non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de sa thèse de quelque manière et sous quelque forme que ce soit pour mettre des exemplaires de cette thèse à la disposition des personnes intéressées.

L'auteur conserve la propriété du droit d'auteur qui protège sa thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

Canada

APPROVAL

Name: Michael Douglas Dorken

Degree: Master of Arts (Psychology)

Title of Thesis: An Episodic-processing Account of Implicit Grammar

Learning: Preservation of the Functional Operations

Performed on Particular Aspects of Item Structure.

Examining Committee:

Chair: Dr. Roger Buehler

_____
Dr. Bruce Whittlesea
Senior Supervisor
Associate Professor

_____
Dr. A. Roger Blackman
Associate Professor

_____
External Examiner
Dr. Peter Graf
Department of Psychology
University of British Columbia

Date Approved: _Aug 20/92_____

PARTIAL COPYRIGHT LICENSE

I hereby grant to Simon Fraser University the right to lend my thesis, project or extended essay (the title of which is shown below) to users of the Simon Fraser University Library, and to make partial or single copies only for such users or in response to a request from the library of any other university, or other educational institution, on its own behalf or for one of its users. I further agree that permission for multiple copying of this work for scholarly purposes may be granted by me or the Dean of Graduate Studies. It is understood that copying or publication of this work for financial gain shall not be allowed without my written permission.

Title of Thesis/Project/Extended Essay

An Episodic-Processing Account of Implicit Grammar Learning:

Preservation of the Functional Operations Performed on

Particular Aspects of Item Structure

Author:

(signature)

Michael Douglas Dorken
(name)

OCT 13/92
(date)

Abstract

People can become sensitive to the structural regularities of an artificial grammar without awareness. After experiencing a set of representative exemplars people are able to classify novel exemplars with above chance accuracy. Numerous accounts of this form of implicit learning emphasize the automatic abstraction of structure. However, because these accounts ignore the functional nature of structural processing, they may be insufficient to explain the variability found in performance. A new approach to implicit learning, the episodic-processing account, emphasizes that different purposes for encountering exemplars of a grammar will cause critically different operations to be performed on particular aspects of stimulus structure. In the experiments conducted, manipulation of the demands of the task in which exemplars were incidentally encountered and encoded was shown to radically modify the nature of the knowledge base acquired in implicit learning. Coding of particular properties of exemplars, under the guidance of the different purposes given to incidentally encounter items, sponsored: 1) distributed representations of the deep structural features of individual items; 2) representations of the surface structural features of the entire set; 3) representations of the deep structural features of the entire set, and; 4) possibly, context-specific representations of the surface structural features of individual items. It is concluded that memory does not automatically abstract structure; instead, memory performs and preserves whatever operations are functional in satisfying current demands.

Acknowledgements

Table of Contents

Table of Contents (continued)

List of Tables

List of Figures

Reber (1989), in a review of over twenty years of experimentation in his laboratory, has put forward a model of human cognition endowing memory with the ability to implicitly abstract the rule-governed regularities embodied in a set of instances. This model has received extended criticism (e.g., Brody, 1989; Brooks, 1978; Dulany, Carlson & Dewey, 1984; Perruchet & Pacteau, 1990; Servan-Schreiber & Anderson, 1990; Vokey & Brooks, 1992), and has generated interest in the incidental learning of complex structural regularities.

The primary interest in implicit learning has centered around the artificial grammar learning paradigm. In this paradigm, subjects are found to be sensitive to the rule-governed regularities of novel grammatical test items after simply memorizing a set of training items generated from the same artificial grammar. The nature of the learning process involved in artificial grammar learning is intriguing. Subjects gain knowledge fostering sensitivity to the underlying structure of the set while unaware that any such structure exists, and the knowledge they gain, knowledge driving their classification performance, is difficult to verbalize and may be unavailable to consciousness (Dienes, Berry & Broadbent, 1991; Reber, Kassin, Lewis & Cantor, 1980; Reber & Lewis, 1977).

Conflicting theories describing the acquisition, form and utilization of implicit knowledge have been forwarded by numerous researchers (e.g., Mathews, Buss, Stanley, Blanchard-Fields, Cho and Druhan, 1989; Perruchet & Pacteau 1990; Reber, 1989; Servan-Schreiber & Anderson, 1990; Vokey & Brooks, 1992). At the current time, the nature of implicit learning is widely debated. Researchers studying artificial grammar learning must uncover the

nature of the memory representations formed while encountering items generated from an artificial grammar, and how these memory representations, once formed, support grammatical sensitivity.

The Standard Artificial Grammar Learning Paradigm

The standard training procedure in the artificial grammar learning paradigm requires subjects to memorize consonant strings generated from a complex set of rules that can be represented in schematic form (see Figure 1). Consonant strings are generated by following the pathways of the grammar from the initial state (State 1) to any of the three terminal states (States 4, 5 or 6). For example, the item MTTC can be generated from the grammar shown in Figure 1. First the M is picked up between States 1 and 2, then the first T by following the loop around State 5, then the second T by following the same loop a second time, and then the C by following the pathway between States 2 and 4. Typical training involves individually presenting twenty-five strings generated from a grammar three times each for ten seconds.
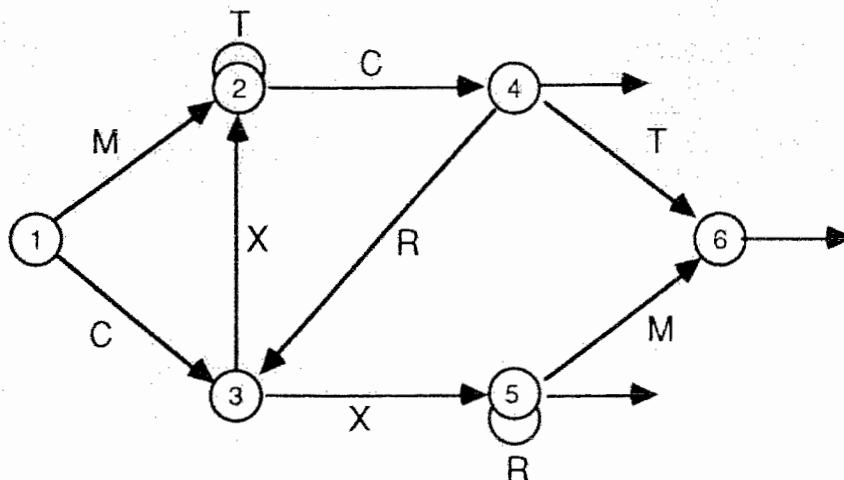
Figure 1. Example of a finite state grammar.

Subjects are not told about the complex set of rules generating the letter strings until they have completed the training task. At test, they are required to categorize novel grammatical items, conforming to the generating rules, and novel nongrammatical items, created to violate at least one of the generating rules. All test items are formed from the same subset of letters (e.g., M, T, V, R, X). Subjects given this standard form of training and test are able to discriminate between grammatical and nongrammatical items in the absence of feedback (e.g., Dienes, Broadbent & Berry, 1991; Dulany et al., 1984; Perruchet & Pacteau, 1990; Reber, 1976).

The Unconscious Structural Abstraction Account of Implicit Learning

Reber (1967, 1969, 1976, 1989) has explained implicit learning by theorizing that memory is capable of veridically representing the complex structure of the environment. Reber has argued that when people are faced with stimuli conforming to underlying structural regularities, such as exemplars conforming to the rules of an artificial grammar, an automatic process abstracts those structural regularities into an internalized representation which, in turn, sponsors structural sensitivity. The structure of an artificial grammar is internalized in a form that, although not identical to the formal Markovian system, is still "deep, abstract, and representative of the structure inherent in the underlying invariance patterns of the stimulus environment" (Reber, 1989, p.226).

Reber (1989) supported the contention that memory's internalized representation directly mirrors the structure underlying the set of exemplars by providing evidence that subjects, given standard training, do not use rules

that misrepresent that underlying structure. Specifically, Reber argued that a knowledge base which misrepresented the underlying general structure of the domain would cause consistent misclassification of items presented more than once at test. In contrast, he argued that a knowledge base that represented the general structure of the domain would result in consistent correct classification of test items presented more than once. Summarizing the results of nine experiments with learning conditions claimed to be neutral, all similar or identical to the standard conditions summarized above, Reber reported that subjects did not make more errors classifying the same items twice than would be expected by estimates of chance. However, these same subjects did correctly classify the same items twice more often than would be expected by estimates of chance. Therefore, Reber concluded that subjects were inducing and utilizing only representative, veridical rules. However, because classification performance was less than perfect, Reber concluded that the rules acquired were incomplete representations of the underlying structure.

The Instance Account of Implicit Learning

The first response to Reber's theory of implicit learning was made by Brooks (1978) who provided an alternative account of implicit learning and concept learning in general. Brooks argued that the acquisition of a concept does not involve abstraction across all experienced instances of a category in order to create summary representations like a prototype or a set of analytic rules, but instead, involves storing all experienced instances. According to the instance view, a novel grammatical stimulus is classified by an analogy to a similar instance. For example, seeing a stimulus (e.g., a dog), causes memory to

automatically access a highly similar instance along with the information that it was a dog, and by analogy the new stimulus is classified as a dog. Brooks argued that the unconscious abstraction of rules across instances of a grammar, although sufficient to explain grammar learning under standard conditions, may be unnecessary; knowledge may be distributed across representations of individual instances.

To demonstrate the power of the instance view, Brooks (1978) performed an experiment using two artificial grammars and a paired-associate learning procedure. In training, exactly half of the letter strings generated from each grammar were paired with an animal associate and half with a city associate. This ensured that an item's grammatical status could not be predicted by the city/animal status of its associate. However, unknown to the subjects, the animal and city associates could be divided into a nonsalient second category that did perfectly predict an item's grammatical membership. Specifically, all letter strings generated from one grammar were paired with an animal or city from the New World, and all letter strings generated from the other grammar were paired with an animal or city from the Old World. Therefore, the nonsalient New World/Old World status of letter string's associate was perfectly predictive of its grammatical membership. In training, subjects memorized the letter strings and their associates. They were not aware that they were in a grammar learning experiment. With respect to the associates, subjects reported only being aware of the nonpredictive city/animal category and not the critical New World/Old World category. Therefore, the subjects were not aware, until informed just before test, that the New World/Old World status of a letter string's associate

was perfectly predictive of its grammatical status. Nonetheless, subjects were successful in sorting novel items into the two grammatical categories.

Because subjects, while learning, were unaware of the critical New World/Old World associate category, Brooks (1978) argued that subjects were not simultaneously abstracting the two different grammatical structures. In support of this view, another group of subjects performed the same acquisition task with the New World and Old World associates randomly assigned to the training items. These subjects experienced the same grammatical items but, contrary to the predictions of an automatic abstraction account, had no ability to discriminate between the two grammars. Therefore, subjects could only discriminate between the two grammars when they could utilize the critical New World/Old World associate category at test. Brooks concluded that successful categorization of novel grammatical items in this experiment must be based upon the use of analogies to similar grammatical items stored in memory with their associates. For example, the test string MRMRV might seem similar to the training string MRRMRV which was associated with Vancouver. Because Vancouver is a New World city, the subject might decide that the test string is also from the New World category. Although the utilization of instances appears to be very explicit in this case, under other circumstances the influence of instances on categorization can be entirely implicit (Allen & Brooks, 1991; Jacoby & Brooks, 1984; Vokey & Brooks, 1992).

The strength of an instance account to explain grammatical sensitivity is further enhanced if, as Vokey and Brooks (1992) suggested, it is assumed that classification depends not simply on the most similar instance as suggested by

Brooks (1978), but rather, on the simultaneous similarity of the current stimulus to multiple known instances retrieved in parallel (Estes, 1986; Hintzman, 1986; Nosofsky, 1988; Whittlesea, 1987). This multiple item retrieval would allow wide generalization to novel grammatical items without necessitating the unconscious abstraction of a set of rules. In fact, this form of on-line computation, involving the simultaneous averaging of stored episodes, produces sensitivity to many of the statistical properties of a domain. Moreover, this form of on-line computation can give identical results to the pre-compution performed by models generating summary representations such as prototypes or rule-sets (Estes, 1986). Therefore, either pre-computation and on-line computation could be responsible for grammatical sensitivity under standard training conditions.

Instances Versus Structural Abstraction

In reply to Brooks (1978), Reber and Allen (1978) agreed that under some circumstances, such as paired-associate learning, the memory base supporting grammatical decisions is in the form of specific instances, but under other circumstances, such as learning under an observation procedure, the memory base is an abstract representation of the rules of the grammar. Reber and Allen argued that memory deploys either a rule abstraction procedure or a more concrete instance-based procedure under the guidance of "the type of material to be learned, the way it is presented, one's expectations about the task, and one's previous success with these procedures" (p. 219). Reber (1989) added that in the pure implicit acquisition mode, in which no overt request or experimental

manipulation causes the use of elaborative operations, the default procedure is abstraction.

Reber and Allen (1978) argued that the differences in recognition and classification performance between subjects given paired-associate training and subjects given observation training supported their dual-knowledge model. Observation training required the subjects to simply pay attention to the items presented. Reber and Allen suggested that instance-based memory representations, when compared to abstraction-based representations, should support greater item-knowledge, as measured by item recognition, but less generalization to novel items, as measured by classification accuracy. In Reber and Allen's study, paired-associate learning resulted in reliably better recognition of old items and reliably poorer generalization to novel grammatical items when compared to observation learning. Therefore, they concluded that paired-associate learning supports the development of a knowledge base of instances and observation learning supports the development of an abstract representation of the rules of the grammar. This conclusion was also supported by the fact that subjects given paired-associate training were ten times as likely to justify their classification decisions by using the words "it reminds me of", which was hypothesized to be a form of justification that should result from the utilization of instance-based knowledge.

This variable item-knowledge and generalization, under control of the induction task, can be explained by the Reber and Allen (1978) model postulating two modes of learning. However, it could also be explained by a form of exemplar model postulating only one mode of learning. Vokey and Brooks (in

an unpublished version of their 1992 article, as cited by McAndrews and
Moscovitch, 1985) suggested that an exemplar model tightly linking item
individuation and breadth of generalization could explain Reber and Allen's
results.

The account proposed by Vokey and Brooks suggested that encoding
manipulations could affect the degree of item individuation, and if items are very
well individuated, they would not support wide generalization to novel items.
However, these highly individuated items would be easily recognized, indicating
good item-knowledge. With respect to the two tasks used by Reber and Allen
(1978), this exemplar model would predict that paired-associate training leads
to greater item individuation than does observation. Therefore, this model would
also predict that paired-associate training should support greater item
knowledge as measured by recognition, but less generalization as measured by
classification. This is the result reported by Reber and Allen in support of their
model postulating dual modes of learning.

Therefore, the results provided by Reber and Allen (1978) can be explained
by a model of instance learning alone, or a model including both instance learning
and abstraction. To determine which model better explained artificial grammar
learning, Vokey and Brooks (1992) and McAndrews and Moscovitch (1985)
attempted to test instance models against dual-knowledge models. Both sets of
researchers attempted to test these models by unconfounding specific similarity,
the similarity of one test item to one training item, with grammaticality, the
similarity of one test item with the entire set of training items. This
unconfounding of specific and general similarity provides the experimenter with

the ability to interpret the source of generalization. Any influence of specific similarity, independent of grammaticality, can be interpreted as the effect of a single, highly similar item in memory controlling classification. In contrast, any influence of grammaticality, independent of specific similarity, can be interpreted as the effect of knowledge pooled across items controlling classification. This latter effect could result from an abstract representation of the rules of the grammar, or parallel access to multiple items at retrieval.

The results of Vokey and Brooks' (1992) experiments demonstrated that both the specific similarity and the grammaticality of the stimuli are determinants of classification performance. However, under various encoding conditions introduced to influence the degree of item individuation, ranging from standard memorization to mnemonic training, variation in the influence of the specific similarity of the stimuli did not result in a compensatory variation in the influence of the grammaticality of the stimuli. Instead, across encoding manipulations, there was no clear tradeoff between the influence of the specific similarity and the grammaticality of the test stimuli.

These results indicated that there is a complex mechanism governing the variable utilization of the different sources of information available to support classification. This complex mechanism is not part of either the dual-knowledge or instance positions. The difficulty facing both types of model centers around their inability to explain why the variation in the effect of specific similarity did not result in a compensatory variation in the effect of grammaticality. The dual-knowledge position (Reber & Allen, 1978) cannot explain this result because it suggested that any training which encourages item individuation, such as

mnemonic training, should cause a switch from the default abstraction mode to an instance learning mode. This switch to the instance learning mode should cause subjects, at test, to rely upon the specific similarity of the test item to a well learned instance (Reber & Allen, 1978). Importantly, this switch should always be accompanied by a smaller reliance upon the grammaticality of items which is hypothesized to be a result of the alternative, abstraction mode of learning. However, across encoding manipulations in the experiments reported by Vokey and Brooks (1992), there was no clear tradeoff between two neatly defined modes of learning controlled by the requirements of the training task.

This same result causes problems for the modified instance account of grammar learning proposed by Vokey and Brooks (as cited by McAndrews & Moscovitch, 1985) which proposed that an increase in item individuation causes a decrease in the breadth of generalization to novel test items. It does not appear that an increase in item individuation necessarily decreases the breadth of generalization to novel grammatical items. Therefore, neither position can easily handle the variable impact of different encoding instructions on item knowledge and the breadth of generalization to novel grammatical items.

Thus, the dual-knowledge position (Reber & Allen, 1978) and the instance account (Brooks, 1978; Vokey & Brooks, 1992) have led to experiments demonstrating the variable reliance of classification performance on different forms of information, but these two theories have been unable to clearly disentangle the mechanisms underlying this variation.

The Fragmentary Knowledge Account of Implicit Learning

In addition to the instance-based and abstraction-based representations, other possible forms of representation have been hypothesized to support sensitivity to the structure of an artificial grammar. With respect to the complexity of the structure necessary to support grammatical sensitivity, Perruchet and Pacteau (1990) have provided the simplest account of grammatical knowledge. They argued that fragmentary knowledge of legal side-by-side letter pairs (bigrams) extracted from grammatical training items, in addition to knowledge of the acceptable initial and final letters in grammatical items, is sufficient to explain the results obtained in standard tests of grammatical sensitivity.

To demonstrate this point, Perruchet and Pacteau conducted an experiment in which subjects memorized legal bigrams whose frequency of presentation matched their frequency of occurrence in the entire set of grammatical training strings. They demonstrated that in a standard test of grammatical sensitivity these subjects performed as well as subjects who had memorized whole grammatical training items. Moreover, the results of another test measuring recognition of legal and illegal bigrams following standard item memorization were included in a simulation programmed to make a judgment of "nongrammatical" for any test string containing at least one bigram not recognized by individual subjects. This simulation predicted the actual classification performance of subjects in a standard grammar learning condition. They concluded that under standard grammar learning conditions subjects

memorizing whole training items may only gain knowledge of permissible bigrams.

Because the location of legal bigrams within items is constrained (e.g., MT might occur exclusively in the first and second position of legal strings), it is possible that subjects may learn not only bigrams but also their positional dependencies. To test this possibility, Perruchet and Pacteau (1990) analyzed their data in terms of nongrammatical test items that violated the rules of item construction by including illegal bigrams and those violating the rules by including legal bigrams in illegal positions. The results of their analysis demonstrated that in the usual grammar learning paradigm a small amount of the grammatical sensitivity may stem from correct rejection of nongrammatical items with permissible bigrams in impermissible positions; that is, subjects were sensitive to the positional dependencies of legal bigrams. However, because this effect was small Perruchet and Pacteau down-played the importance of positional information under standard grammar learning conditions.

The potential sensitivity of subjects to the positional dependence of bigrams was further examined by Dienes et al. (1991). Using a sequential letter dependencies (SLD) test, they were able to directly test subjects' sensitivity to the positional dependencies of bigrams. This test required subjects to state which letters were allowed to follow a stem of one to five letters extracted from the beginning of a grammatical item. They demonstrated that subjects, after receiving standard memorization training, were sensitive to the positional dependence of bigrams. In contrast, sensitivity to the positional dependence of bigrams was absent when subjects were required to simultaneously generate

random numbers and study the training items. Therefore, it appeared that knowledge of both valid bigrams and their positional dependencies were learned under standard conditions; however, it also appeared that the knowledge of positional dependencies developed only after more extensive practice than is necessary to simply learn the bigrams independent of their positions.

Adding to this line of research, Servan-Schreiber and Anderson (1990) demonstrated that when subjects received extended practice memorizing the letter strings (to a criterion of two successive correct reproductions of each set of five exemplars), their performance could be simulated by a model postulating that subjects learned small chunks (bigrams and trigrams) of each letter string and then integrated these small chunks into a hierarchy of more encompassing chunks. In support of the contention that chunking is an important process involved in artificial grammar learning, Servan-Schreiber and Anderson induced subjects to chunk the letter strings in a way consistent with the specific organization imposed by the training presentation. As a result of this training manipulation, subjects were least sensitive to violations in letter ordering that occurred between the chunks of specific letter sequences they had originally encoded. Subjects were most sensitive to the violations that occurred within the chunks of specific letter sequences that they had originally encoded. Importantly, they also demonstrated that subjects not induced to chunk the items by some overt manipulation also showed this same dependence upon chunk preservation.

Therefore, it would appear that under standard memorization training with grammatical items knowledge of bigrams and possibly trigrams develops along with some knowledge of the positional dependencies of these groupings, and then if

more practice memorizing items is allowed, bigrams and trigrams are integrated into even larger groupings forming representations of whole items. It is also possible that the entire set of small, intermediate and item sized "chunks" are represented in an all-encompassing hierarchical network of "chunks" (Servan-Schreiber & Anderson, 1990).

Fragmentary Knowledge Versus Structural Abstraction

On the surface, describing the form of the knowledge acquired in standard artificial grammar learning conditions as bigrams, trigrams and their locations contradicts Reber's (1989) contention that subjects acquired knowledge that can be characterized as deep, abstract, and representative of the structure underlying the stimulus environment. Avoiding this contradiction, Reber (1989; Reber & Lewis, 1977) claimed that the knowledge acquired in implicit learning is in the form of permissible bigrams, but that this knowledge of bigrams reflects the deep structural characteristics of the grammar. Reber and Lewis claimed that memory first abstracts the bigrams that have the highest relational invariances and then abstracts the less invariant bigrams. This abstraction, supported by sensitivity to the deep structural regularities of the domain, was hypothesized to be independent of the frequency of bigram occurrence in the training set. To support this position and refute any account of grammar learning positing that subjects are simply sensitive to the frequency of bigram occurrence in the training set, Reber and Lewis provided the results of an anagram solution test. This test required subjects to unscramble anagrams, randomly mixed letters from novel grammatical items, into grammatical items. The frequency with which each bigram occurred in the subject's solutions to the

anagram task was then correlated with: a) their frequency of occurrence in the training strings (r = .04), and b) their frequency of occurrence in the entire set of grammatical items (r = .72). From these data, Reber and Lewis concluded that subjects were "learning the overall structural relations that hold between letters and letter pairs and not simply logging frequencies [of letter pairs in the training set]" (p. 347).

However, Perruchet and Pacteau (1990) uncovered several important problems with Reber and Lewis' (1977) study. Specifically, Perruchet and Pacteau suggested that the impact of three methodological factors in the study conducted by Reber and Lewis could explain the fact that the frequency of occurrence of bigrams in the anagram solutions most closely matched their frequency of occurrence in the entire set of grammatical items, and not their frequency of occurrence in the training set. First, Perruchet and Pacteau pointed out that the variance in the distribution of bigram frequency of the items selected for training was approximately one ninth of the corresponding variance calculated for the entire set. This comparatively low variability in the training bigram frequency distribution could lead to the result reported by Reber and Lewis because it makes it difficult to obtain as high a correlation with the bigram frequency of the whole set.

Second, Perruchet and Pacteau (1990) pointed out that the individual letters used to create the anagrams were more representative of the frequency of individual letters in the entire set than the frequency of individual letters in the training set. This restriction would naturally cause the bigrams, created by

combining this restricted set of individual letters, to match the bigrams in the entire set more closely than the bigrams in the training set.

Third, Perruchet and Pacteau (1990) pointed out that the bigrams that were underrepresented in the original training list were highly salient "doublets" (VV, XX, TT) and a highly salient abbreviation (TV); however, the bigrams that were overrepresented were non-salient (PV, XS and PX). As a result, a misrepresentative overuse of these highly salient and easily memorized bigrams in the subsequent anagram solution task would reduce the overall correlation between the bigram frequency in the training set and the anagram solution test. Together, these three biases raise questions about the proper interpretation of Reber and Lewis' (1977) results because none of them require more than "the ubiquitous ability to learn pairwise associates" (Perruchet & Pacteau, 1990, p. 271).

In the absence of Reber providing a plausible mechanism to explain how memory gathers bigram knowledge under the guidance of deeper knowledge of the underlying grammatical structure, and how this deeper structure is apprehended in the first place, the arguments made by Perruchet and Pacteau (1990) are compelling. It is far more parsimonious to suggest that subjects in the standard grammar learning paradigm simply learn bigrams and to some extent their positional dependencies.

Fragmentary Knowledge in the Form of Condition-Action Rules

Mathews and his associates (Druhan & Mathews, 1989; Mathews, 1990; Mathews, 1991; Mathews et al., 1989; Roussel, Mathews & Druhan, 1990) have argued that bigrams and trigrams and their positional dependencies may be

internally represented in the form of specific condition-action rules. To develop this position, Druhan and Mathews (1989) have created a model, THIYOS (THe Ideal YOked Subject), based upon the classifier systems proposed by Holland, Holyoak, Nisbett and Thagard (1986). THIYOS was originally used by Druhan and Mathews to mimic the performance of yoked subjects given the verbal protocols created by subjects who had performed in the Mathews et al. (1989, Experiment 1) grammar learning experiment. In the condition that was simulated, subjects performed in a string discrimination task in which they had to select one of five exemplars presented as grammatical. The set of exemplars was composed of one grammatical exemplar and four nongrammatical exemplars. Subjects were given feedback after each trial and the correct choice remained on the screen for five seconds. After every ten trials they had to verbalize "teach-aloud instructions" to an "unseen partner" who was also performing the task. These verbal protocols of each subject were given to a yoked subject who performed the same task without feedback. This practice continued for 200 trials per week for three weeks.

THIYOS performed the string discrimination task by first storing the verbal protocols of human subjects modified into specific condition-action rules (for example, if the string begins with SCT then choose it) and then modified again into a specific format amenable to computer programming. These rules competed for control of response selection based upon the parameters of strength, specificity and support from other rules. In this model the strength of a rule was based upon its past success in guiding classification, the specificity of a rule was based on the number of the positions within a string it uniquely specified, and the

support for a rule was based on the number of other rules that were in agreement with the decision it indicated. In a first round of bidding, the rules were compared with the stimuli presented and the strongest rules that match the stimuli were sent on to working memory. Once in working memory, these rules "can be consciously manipulated based on their strength, specificity and support" (Druhan & Mathews, p. 7). This conscious manipulation was modelled by an equation calculating a second set of response bids for every rule. A bid was created for each rule by substituting the specific values of the parameters (strength, specificity and support) attached to each competing rule, weighting them, adding them together and multiplying the result by a constant. The highest bidder governed the response made and, if the system was given feedback, had its strength increased if it led to correct classification or decreased if it led to incorrect classification. Without feedback, this system performed at a level similar to that of human yoked subjects. Moreover, the inclusion of feedback allowed this system to perform at a level similar to that of the original subjects.

Roussel, Mathews and Druhan (1990) modified THIYOS by adding a "forgetting" algorithm which allowed it to retain a subset of features of any item the existing system declared grammatical. These features were coded in the form of a rule that said to select strings with those features. Roussel et al. also added an internally generated feedback algorithm which responded to the structure of the exemplars presented. These modifications allowed the system to learn without the input of the rules generated from subjects' verbalizations and without feedback; in fact, it learned so well it performed well above chance on the two grammar tasks it was tested on. However, Roussel et al. reported that human

subjects in the Mathews et al. (1989) study did not exceed chance on the same tasks. Roussel et al. explained that these differences resulted from human subjects explicitly generating and maintaining hypotheses of little or no validity while THIYOS did not.

Mathews and his associates are strict believers in the psychological realism of their computational model: they assume that their computational model (THIYOS) is the psychological model of implicit grammar learning (cf. Mathews, 1991). This psychological model involves human memory re-coding the features (sequences of letters and their locations) of items deemed to be grammatical into specific condition-action rules that could unconsciously compete to create optimal relative strengths for each competing rule (Mathews, 1991). The rules themselves are variably available to consciousness depending upon their strength, but the mechanisms creating these rules, selecting them for consciousness and optimizing their strength, are not. Once in consciousness rules are manipulated according to an equation based upon their strength, specificity and support.[1] In summary, their belief that a "good theory of implicit learning is close at hand" (Mathews, 1991, p. 118) revolves around the

---

[1] It is difficult to interpret what Mathews and his associates meant when they stated that their second bidding process is "conscious". Clearly, subjects in their experiment were never aware of substituting numbers into a complicated equation. Mathews and associates appear to assume that working memory and consciousness are very tightly linked. However, if this bidding process occurred in working memory, it would not have to be available to consciousness. Only the end result, in the form of a rule or a decision, would have to be made consciously available.

notion that "a large part of the nonverbalized, tacit knowledge acquired about an artificial grammar appears to be the optimal relative strengths of competing rules resulting from the nonconscious rule-tuning" (Druhan & Mathews, 1989, p. 8).

Contrary to the statements made by Mathews and his associates, the limited success demonstrated by THIYOS does not indicate, by any means, that there is a psychological process occurring analogous to the processes performed by THIYOS. Although such a representation of knowledge successfully simulated experimental results to some degree, it does not mean that it is relevant to any actual psychological process (cf., Kolers & Smythe, 1984, for detailed arguments concerning this issue). In fact, any number of alternative representation systems can explain a small set of results from psychological experiments (Anderson, 1978, 1979). Therefore, when a number of very different accounts have the ability to explain a small set of behavioral results, the ability to explain these results becomes an insufficient reason to prefer any one of those accounts. With respect to implicit learning, Servan-Schreiber and Anderson's Competitive Chunking Model (1990), the PDP model created by Kushner, Cleermans and Reber (1991) and THIYOS, all with radically different processing and representation systems, can all adequately explain a specific set of behavioral results. Because such a diverse set of models can explain the results of artificial grammar learning studies, the only true support for any model must come from the results of psychological experiments, in which each model makes numerous specific and unique predictions (cf., Anderson, 1978, 1979, for arguments that even this may not be sufficient). The validity of THIYOS as a psychological model

of implicit learning will be discussed briefly with respect to the psychological data it is based upon, and later with respect to the studies presented in this thesis.

THIYOS is an attempt to simulate a set of results from psychological experiments and, as yet, has made no actual predictions concerning artificial grammar learning. It is a computational model with possible relevance as a psychological model. This distinction between computational and psychological models is an important one. For example, there are computational models successful at playing chess at the Grand Master Level; however, such models may or may not have relevance to psychologists with respect to their ability to act as models of how expert humans play chess (Kolers & Smythe, 1984). On the whole, such models seem to tell us little about human chess expertise. However, there are other computational models such as Servan-Schreiber and Anderson's Competitive Chunking Model (1990) which can also inform psychologists about human performance by acting as psychological models. The Competitive Chunking Model is a successful model within a limited domain because it is backed by behavioral evidence that supports its main operating assumption. This assumption states that people chunk information, and in the domain of artificial grammar this has very specific behavioral consequences. This model, in fact, predicted and modeled the findings presented by Servan-Schreiber and Anderson (1990) in a more specific way than any other psychological model. To be similar in its ability to act as a valid psychological model THIYOS must also lead to interesting psychological predictions which it is capable of modelling (Estes, 1986; Kolers & Smythe, 1984).

One possible form of confirmation for THIYOS would be finding behavioral support for its representational assumptions. Mathews and his associates made the suggestion that the verbal reports given by subjects were a valid index of the form of the representation of knowledge implicitly supporting decisions (Mathews et al., 1989). In fact, they justified the use of rules in their model on the nature of the verbal reports of the subjects in a number of experiments conducted by Mathews et al. (1989). If it could be demonstrated that the verbalized knowledge given by subjects following the teach-aloud procedure was in the form of condition-action rules, which were hypothesized to be consciously available, then there would be at least some support for the psychological relevance of THIYOS. However, even this would be exceptionally weak support: verbal statements in the form of condition-action rules could be a product of the "teach aloud procedure", rather than a reflection of true "implicit" knowledge. Statements in the form of "if you see X then choose that string" are the most valid of any possible verbal information that can be explicitly generated by subjects and then passed on in verbal form. This in no way means that the knowledge fostering implicitly-driven decisions is in a corresponding form. Almost any information about the features of items could be reformulated into a condition-action rule by either the experimenter, a computer induction routine, or the subjects themselves attempting to come up with something to say into a tape recorder. Thus, even if the subjects did speak in perfect condition-action rules there is no reason to believe that this was the form of their implicit knowledge. To underscore this point, Allen and Brooks (1991) demonstrated that even when subjects were given a very simple, perfectly predictive rule in verbal form to

categorize simple stimuli their performance relied on prior episodes rather than rule automatization.

This argument against the validity of verbal reports as an index of implicit knowledge, although powerful, is perhaps unnecessary because the verbal reports given by subjects were not in the form of condition-action rules as they should have been if they are a direct reflection of implicitly-held knowledge. In fact, Druhan and Mathews (1989) thanked their scorer for "translating scores of verbal transcripts into regular expressions that could be parsed into classifiers" (p. 9). Therefore, there is no evidence that human subjects are using rules in the form necessitated by THIYOS even though some of those rules, as argued by Mathews and his associates, should be directly available to consciousness.

The most critical problem that may totally undercut the validity of THIYOS as a psychological model of implicit learning is the fact that it was based upon and simulated situations in which subjects may have learned explicitly. The most successful simulation was conducted by Druhan and Mathews (1989) who used the data from two human subjects who performed a task, described above, which would appear to orient the subjects toward active, explicit hypothesis testing. Subjects made their choice of the most grammatical item out of a set of five items, and then were given feedback in the form of the correct answer. Therefore, subjects were given every opportunity to form and test explicit rules about the correct stimuli on each trial and, in fact, they were encouraged to do so by the fact they had to verbalize their knowledge every ten trials.

Surprisingly, Roussel et al. (1990) argued that it was their no-feedback condition, similar to the more usual grammar learning conditions, that caused subjects to suddenly use "explicit processes". These explicit processes were hypothesized to cause the generation of invalid rules, thereby interfering with the usual implicit processes that occurred with feedback. Their only evidence supporting this claim was that explicit processing, invoked by instructions to search for rules, has been reported by Reber (1976) to cause poor performance on grammar learning tasks. However, this seems rather weak evidence considering that this same instructional manipulation was a part of the Mathews et al. (1989) study and had no reliable effect on learning (cf. Dienes et al., 1990; Dulany et al., 1984; Perruchet & Pacteau, 1990, for other failed attempts at replicating Reber, 1976). Therefore, Roussel et al. must argue that giving subjects instructions to explicitly search for rules does not lead to explicit processing, but removing feedback does.

A much simpler account of the Mathews et al. (1989) study is that the nature of the learning task, which included extended practice and verbal justification every ten trials, caused subjects to use explicit rule generating strategies in the absence or presence of feedback or instructions to proceed explicitly, and that this explicit strategy was successful only when the feedback supported the selection of appropriate rules. Therefore, subjects in the task being modelled by Mathews and his associates were very likely using explicit knowledge.

It does not appear that there is any evidence suggesting the model presented by Mathews and his associates is relevant to the psychological phenomenon of implicit learning. The formation of condition-action rules and their optimization

is a well defined "in principle" method by which implicit learning could occur. Although bigrams and their positional dependencies can be described as condition-action rules, there is no evidence that there is a psychological mechanism that performs this transformation.

The Importance of Transfer Across Letter Sets

Although Reber's theory of implicit learning has been placed in question by the numerous studies on the standard grammar learning conditions, a line of evidence still exists suggesting that there is more to grammar learning than simply coding items or parts of items. This evidence comes from studies in which subjects classify test items instantiated upon a letter-set different from the one used to instantiate the training items. For example, subjects might study grammatical items instantiated upon the letter set {M, T, V, R, X} and then at test classify grammatical and nongrammatical items that have been translated into the letter set {Z, J, F, K, L}. For example, by consistently replacing every M with a Z, every T with a J, every V with a F and every R with a K, the item MTRVM becomes ZJKFZ. Experiments by Reber (1969) and Mathews et al. (1989) have demonstrated that the knowledge gained from memorizing grammatical items instantiated upon one letter set can support grammatical sensitivity to items generated from the same grammar but instantiated upon a different letter set.

Reber (1989) claimed that this changed letter-set transfer cannot rely upon knowledge that is based on the superficial physical form of the stimuli. Instead, it must rely upon "knowledge of the deeper, more abstract relations that can, in principle, be said to underlie them" (p. 225). Reber's full argument can be stated as follows: 1) strong transfer across letter sets demonstrates that the

knowledge supporting grammatical sensitivity is not tied to the specific letter-set on which the training items are instantiated; 2) a knowledge base consisting of only learned instances or parts of instances would be tied to the letter set used to instantiate the instances, and therefore; 3) the knowledge base that leads to strong changed letter-set transfer cannot consist of learned instances or parts of instances. Reber concluded by postulating that the only possible form of knowledge not tied to the letter set used to instantiate items, and therefore the form of knowledge that must govern changed letter-set transfer, is a mental representation of the abstract rule-governed regularities that define the structure of the domain.

Mathews and his associates (Mathews, 1990, 1991; Mathews et al., 1989) are in general agreement with Reber's claim. Mathews (1991) reported that THIYOS, the rule induction system, could perform above chance on changed letter-set transfer tests if the occurrence of repeated letters or runs and their spatial locations were transformed into specially created condition-action rules.

It is generally accepted that Reber (1989) was correct in stating that a knowledge base consisting of parts of items (e.g., bigram knowledge proposed by Perruchet & Pacteau, 1990) cannot account for changed letter set transfer. For example, Mathews (1990) argued that sensitivity to the grammatical status of items across letter sets necessarily involves location information. In addition, Servan-Schreiber and Anderson (1990) conceded that their model of grammar learning, relying on chunking of particular letter sequences, could not explain changed letter-set transfer.

However, Reber's proposition that a knowledge base consisting of whole instances cannot explain this form of transfer has been contested. Brooks and Vokey (1991) proposed a mechanism by which changed letter-set transfer could be supported by similarity to a set of learned instances. Introducing the term abstract analogy (Gentner, 1983) to the artificial grammar paradigm, Brooks and Vokey pointed out that there can be an abstract similarity between items instantiated upon different letter sets. For example, MXRVVVM can be considered similar to BDHCCCB because they both begin and end with the same letter and have a repeated letter triplet in the same position. They claimed that whole instances are stored in memory and classification is governed by a similarity mechanism; if a test item is similar enough to an item stored in memory on an abstract, relational dimension then the person will call the item grammatical.

In an attempt to test their position, Brooks and Vokey (1991) used materials in which the similarity of test items to specific training items was unconfounded from the grammaticality of the test items. They concluded from the results of their study that changed letter-set transfer, similar to same letter-set transfer, was supported by both specific similarity to individual training items and by general similarity to the whole set of training items. The latter, according to Brooks and Vokey (1991), could be explained either in terms of an abstract representation of the grammar or the abstract similarity of each test instance to multiple training instances stored in memory.

In summary, strong arguments based on demonstrations of changed letter-set transfer have not settled the debate over what form of representation supports grammatical sensitivity.

A Processing Framework of Memory

The accounts of implicit learning that have been reviewed all differ in the kind of structure that they require the mind to acquire while learning. Although some accounts stress that the mind is sensitive to bigrams, some whole instances, and still some others abstract rules, they all stress a particular form of structure and how this structure, once transferred to memory, can explain grammatical discrimination. There is, however, another approach to the nature of learning. This approach, exemplified by memory for procedures (e.g., Kolers, 1979; Kolers & Roediger, 1984), focuses on the specific operations that the mind performs in order to deal with stimulus structure. This view stresses the critical importance these operations have in determining the knowledge base acquired, and the highly specific ability of that knowledge base, once acquired, to support later performance.

Kolers' approach, focussing on the operations performed in dealing with stimuli, combined with concepts such as encoding variability (Martin, 1971), encoding specificity (Tulving & Thompson, 1973) and transfer-appropriate processing (Morris, Bransford & Franks, 1977), is the backbone of a theory of memory focussing on the processing performed when dealing with stimuli (e.g., Jacoby, 1983; Roediger, Weldon & Challis, 1989; Whittlesea, 1987; Whittlesea & Brooks, 1988; Whittlesea & Cantwell, 1987). This general processing framework, primarily based upon the ideas put forward by researchers in the 1970's (e.g., Craik & Tulving, 1975) to explain performance on explicit tests of memory, has been extended by current researchers to explain performance on implicit tests of memory and dissociations between implicit and

explicit tests of memory (e.g., Jacoby, 1983; Masson & MacLeod, 1992; Roediger, 1990; Roediger, Weldon & Challis, 1989).

In addition, by clearly defining the processing framework and its implications, Whittlesea and his associates have been able to demonstrate its breadth and scope, using it to explain phenomena as widely divergent as the word superiority effect (Whittlesea & Brooks, 1988), pseudoword perception (Whittlesea & Cantwell, 1987), semantic priming (Whittlesea & Jacoby, 1990), illusions of memory driven by perceptual and conceptual fluency (Whittlesea, 1992; Whittlesea, Jacoby & Girard, 1990) and sensitivity to typicality in clustered categories (Whittlesea, 1987).

The Episodic-Processing Account

The extended processing framework put forward by Whittlesea and his associates has been named the "episodic-processing account" to reflect the fact that it "emphasizes the processing conducted within particular experiences as the primary explanatory mechanism of memory, and that the processing conducted depends on the particular demands and affordances of the encoding episode" (Whittlesea & Dorken, 1992, p. 12). Under this view, memory does not preserve stimulus structures, but instead preserves the products of operations performed in dealing with particular aspects of those stimulus structures. The representation formed in any given experience with a stimulus is driven by the perceived purpose of the encounter, which selects the processes to be performed in an interaction with the structure of the stimulus and the person's past processing experience. The structure of the stimulus and past experience,

together, guide and limit the type of processing that may be applied to the structure of a stimulus, and the efficacy of that processing, once applied.

Whittlesea (1987) conducted a set of experiments that demonstrated the power of the episodic-processing account's ability to explain variable sensitivity to structural regularities. Whittlesea demonstrated that systematic variation in the demands of an encoding task had a large and consistent impact on the sensitivity to the set-wise structure, in this case typicality in a clustered category space, underlying a set of instances. Given the same training items carrying the same objective structure, but different encoding tasks, subjects were sensitive to the typicality of a test item, or alternatively, the similarity of test item to a particular training item. This alternating sensitivity was predicted by an account postulating that stimuli are encoded as experienced and retrieved in parallel.

Specifically, Whittlesea (1987) concluded that this alternating structural sensitivity was a result of variable encoding of the training stimuli controlled by the learner's intentions, the demands of the encoding task and the affordances of the structure being processed. Encoding was found to be variable across encoding conditions. However, within each encoding condition specific forms of stimulus encoding resulted in specific representations of particular processing experiences which served to selectively support later processing at test. This selective support could not depend exclusively upon the similarity of current and past objective structure which remained constant across encoding conditions, but rather, depended upon the similarity of the current and past objective structure as it was processed. Therefore, sensitivity to general structure, in this case

typicality, of a set of items appeared to be supported by the preservation of distributed representations of particular experiences with the structure of individual items (Whittlesea, 1987; Whittlesea, Brooks & Westcott, 1992). It is this same form of memory which is hypothesized to underlie sensitivity to the grammatical status of items generated from a grammar.

The Episodic-Processing Account and Implicit Learning

According to the episodic-processing account, encoding variability and transfer-appropriate processing are important determinants of performance in the artificial grammar learning paradigm. Memory can process structure in many ways, with each having a particular influence on later performance. This influence on later performance is governed by the similarity between the operations performed in gaining knowledge about training stimuli and the operations necessary to efficiently apply that knowledge to the test task.

The Episodic Processing Account Versus Structural Accounts

According to the episodic-processing account, the issue of whether memory acquires bigrams, whole instances or set-wise structural abstractions is, in fact, a secondary issue. Memory is capable of incorporating all these forms of structure, but only under the guidance of specific operations performed to deal with these forms of structure. However, the principle of encoding variability states that even once incorporated, memory does not hold these forms of structure in some standard format, but rather, as they were processed. According to the episodic-processing account, for example, an instance is stored in memory in the way it was originally processed for a particular purpose. The use of mnemonic training by Vokey and Brooks (1992) to create "highly

individuated" instances may not cause memory to code the instances per se, but rather, may cause memory to code the operations performed in generating the mnemonics, repeating them and making them meaningful. According to the episodic-processing view, it is the highly distinctive products of the operations performed upon instances, rather than the instances themselves, that are preserved by memory and guide future performance. The episodic-processing account is, however, similar to the instance account in that they both assume that knowledge is acquired by storing experiences of individual events.

Because the theories of implicit learning to date have been primarily concerned with structure, the influence of encoding variability on the form of knowledge gained in artificial grammar learning has not been well examined. In fact, most studies have used standard memorization training instructions under the assumption that such an induction task is neutral with respect to the stance the subject takes toward the stimuli (Reber, 1989). Specifically, Reber argued that memorization instructions do not encourage or direct subjects to perform any special operations upon the stimuli that would discourage automatic abstraction. However, subjects given standard memorization training instructions, depending upon their interpretation of the appropriate way to proceed, may attempt to: continually repeat the sequence of consonants from left to right; pronounce the item by adding their own vowel sounds; chunk the item into more meaningful subunits (e.g., MR and TV); reorder the consonants into a more personally meaningful order; make mnemonic phrases out of the sequence of consonants; or memorize the location and identity of repeating consonants. These are all encoding strategies which, according to the episodic-processing

view, should have an influence on the representation of grammatical knowledge. However, their influence has not been measured because variables affecting encoding, except in a few cases, have not been independently manipulated.

In contrast to most studies of implicit learning, a study conducted by Servan-Schreiber and Anderson (1990) manipulated the form of encoding required in training. In fact, this study demonstrated a modulation in the form of grammatical knowledge under the control of encoding demands. Servan-Schreiber and Anderson induced subjects to chunk training strings in a way consistent with the specific organization imposed by the visuo-spatial characteristics of the training item presentation. As a result, at test, subjects were less sensitive to violations in letter ordering that occurred between the sub-sequences they had been induced to code, than to violations that occurred within those sub-sequences.

This processing effect could not be explained by an instance account, a fragmentary knowledge account or a structural abstraction account, because they deal exclusively with the structure of the stimuli and not the way that structure is processed. However, the episodic-processing account is capable of explaining this effect. The episodic-processing account suggests that the overt or incidental demands of the encoding task, in this case the visuo-spatial characteristic of item presentation, modulate the operations performed on the stimuli, causing a particular organization of the structure of individual items within memory. This organization of structure, as processed, is then preserved by memory. The structure of test items should then be organized according to the method preserved by memory. Therefore, if a test item should be judged nongrammatical

only if its sub-sequences, as encoded, are dissimilar to the sub-sequences that subjects were induced to code in training. That is, according to the episodic-processing account, grammatical sensitivity is dependent upon the similarity of past and present items as organized by the operations performed in dealing with item structure. This was the case in Servan-Schreiber and Anderson's experiment.

The nature of the structural regularities supporting grammatical sensitivity under different circumstances are of critical importance to any theory of implicit grammar learning, including the episodic-processing account. Before outlining the experiments performed in this paper, an explanation of the possible nature of the structural regularities underlying a set of grammatical items will be discussed. Memory's potential sensitivity to these forms of structure will then be discussed in light of the episodic-processing account and the current set of experiments.

The Nature of Structural Regularities

It is obvious that any account of implicit learning must define the nature of the structural regularities underlying grammatical sensitivity. To this end, Reber (1969) argued that the subjects are sensitive to the set of deep structural rules embodied in the formal diagram of a grammar. However, the Markovian system used by the experimenter to generate grammatical stimuli is only one of an infinite number of possible rule sets capable of describing a specific set of grammatical items (Dulany et al., 1984; Servan-Schreiber & Anderson, 1990; Reber, 1989). Therefore, Reber (e.g., 1989) revised his argument, suggesting instead that subjects internalize a set of rules that, although not identical to the

formal grammar as defined by the experimenter, is similar in form. These rules are capable of generating the structural regularities of the set of grammatical items, and therefore, were hypothesized to support grammatical sensitivity under both same and changed letter-set transfer tests. However, this description is only one of the many qualitatively different descriptions of the structural regularities underlying a set of grammatical items. In fact, more rigorous descriptions of the character of the domain are possible.

In a contrasting account of the structural regularities made available by a set of grammatical items, Mathews (1990) suggested that subjects have knowledge of bigrams and their positions, and more abstract knowledge of the runs and repetitions of individual elements within items. Mathews argued that only knowledge of runs and repetitions can support above-chance performance in the new search space created when the letter-set instantiating items is changed. Mathews was pointing out that very different kinds of structural regularities allow discrimination between grammatical and nongrammatical items under same versus changed letter-set conditions. This form of argument, and the implications of Brooks and Vokey's (1991) abstract analogy mechanism, can be extended to explain the possibly critical aspects of item and set structure.

Set-wise structural regularities must, by definition, rely upon the cross-item commonalities in structure found in the set of grammatical items. That is, each and every item in a set of grammatical items has structural characteristics, and the frequency with which each of these characteristics are found to occur across the set determines the set-wise structure of the set. Therefore, before set-wise structural regularities can be understood, the nature of the structure of

individual items must be determined. In fact, Brooks and Vokey's (1991) discussion of grammatical items suggests that there are two critical aspects to item structure. The first form of structure exhibited by an exemplar of an artificial grammar is its **surface structural features**. A surface structural feature is one of the elements, such as a letter, used to instantiate grammatical items. Surface structural features can be incorporated into specific surface structural sequences. For example, the item CXCRXR has among its many surface structural sequences: C, X, CX, RX, CXC, XCR, CRX, CRXR, XCRXR and CXCRXR. These features are defined as surface structural because they are specific to the elements (e.g., letters) chosen to instantiate the items.

In addition, general, or set-wise, surface structural regularities can be described as the frequency with which certain surface features, or sequences of surface structural features, or sequences of surface structural features and their positions occur in the entire set of grammatical items. Information concerning the surface structural regularities of a set of grammatical items can, in principle, support transfer from one set of grammatical items to another set of items generated from the same grammar and instantiated upon the same letter set. Grammatical items based upon the same letter-set share more common sequences of surface structural features with themselves than with nongrammatical items and therefore, this form of information could support grammatical discrimination (Perruchet & Pacteau, 1990). However, this form of information could not support grammatical discrimination across letter-sets because grammatical items based upon different letter-sets do not share surface structural features.

Importantly, the inclusion of different lengths of surface structural sequences (e.g., bigrams or trigrams or whole string sequences) and their positions modulate the form and usefulness of surface structural information. For example, information pertaining to the frequency of each individual surface structural feature in the training set of grammatical items could support grammatical sensitivity if the frequency of occurrence of individual surface structural features in the training set more closely matched their frequency of occurrence in grammatical test items than in nongrammatical test items. However, in most sets of items created by other researchers, the set of nongrammatical items do not differ appreciably from the set of grammatical items in their frequency of occurrence of individual letters. Therefore, information pertaining to the frequency of pairs or triplets of individual surface structural features and their positions may be a more predictive form of information.

To explain the second aspect of item and set structure, it is helpful to first examine the procedure used to generate items from an artificial grammar. Any item generated from an artificial grammar is formed by picking up elements along the pathways of a grammar. For example, the item CXCRXR can be generated from the grammar shown at the top of Figure 2. First the C is picked up between State 1 and 3, then the X between State 3 and 2, the C between 2 and 4, the R between 4 and 3, the X between 3 and 5, and finally, the R by following the loop around State 5.
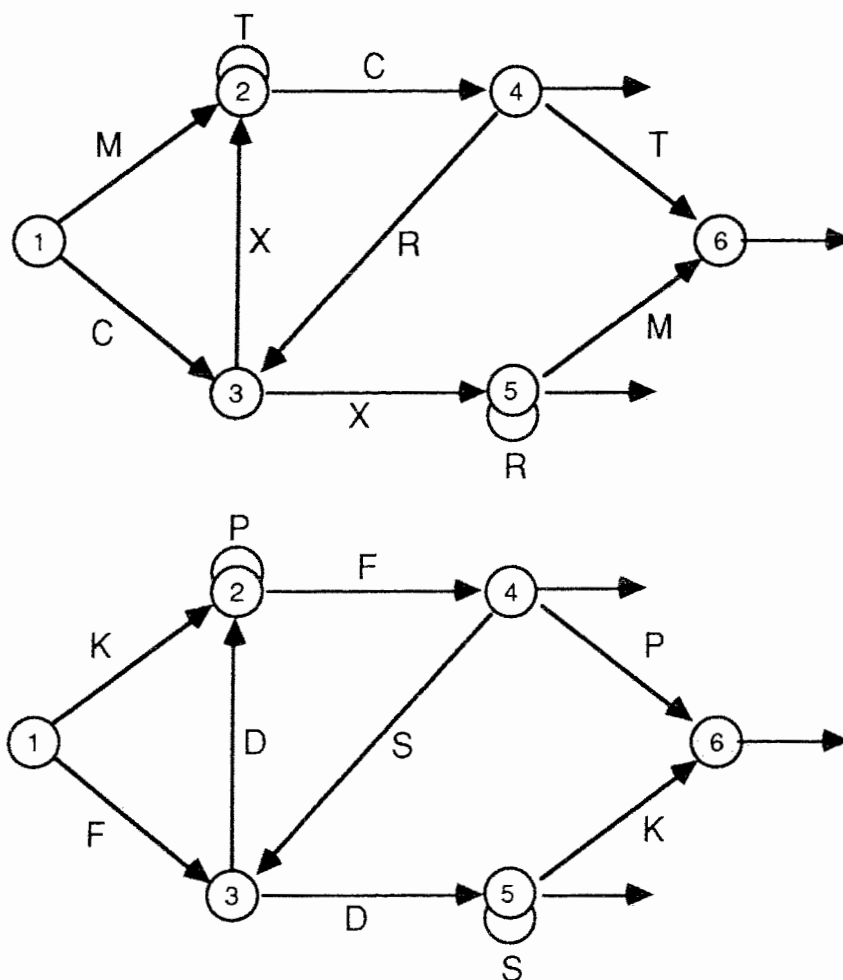
Figure 2. A finite state grammar instantiated with two different letter-sets.

At the bottom of Figure 2 is the same grammar diagram with a different letter-set used to instantiate the grammar. When performing a changed letter-set transfer experiment, the pathways and their linkages are permanent but the letters along the pathways are consistently replaced. In this example, the letter M in the top diagram was consistently replaced in bottom diagram by the letter K, the letter T by P, the letter R by S, the letter C by F and the letter X by D. With this form of letter substitution, the pathways that generated the item CXCRXR in

the top diagram can then generate the item FDFSDS in the bottom diagram. Although CXCRXR and FDFSDS are created from the same pathways of the same grammar, they are not the same item because they are formed from different letters. The only similarity between the two items is that they share the same pattern of element repetition (Brooks & Vokey, 1991). That is, in both CXCRXR and FDFSDS the element in the first position is repeated in the third position, the element in the second position is repeated in the fifth position, and the element in the fourth position is repeated in the sixth position.

The nature of this underlying relationship between any two items generated from the same pathways of a grammar using different elements suggests that there is a second aspect of item structure. Specifically, exemplars of an artificial grammar have **deep structural features**. A deep structural feature is the specific repetition pattern of a single element in any item. For example, a deep structural feature of the item CXCRXR is the occurrence of the element C in the first position and again in the third position. These features are described as deep structural because they are independent of the identity of the elements from which the item is composed. It is the entire set of deep structural features of an item which can define that item as being the same as another item based upon a different letter-set. For example, CXCRXR and FDFSDS have an identical set of deep structural features.

By extension, general, or set-wise, deep structural regularities can be summarized as the frequency with which each deep structural feature occurs across the entire set of grammatical items. This form of structural regularity could be expressed as the frequency with which individual deep structural

features (e.g., repetition of an element in the first and third positions) occur in the entire set of grammatical items, or as the frequency with which conjunctions of deep structural features (e.g., repetition of an element in the first and third positions and repetition of a different element in the second and fourth position) occur in the entire set.

Information concerning the deep structural regularities of the set of grammatical items can, in principle, support grammatical discrimination in both same and different letter-set transfer tests. Grammatical items, based upon any letter set, necessarily share more deep structural features or conjunctions of deep structural features with each other than with nongrammatical items because of the generating principles of an artificial grammar.

This view of set-wise structure is in contrast to the view proposed by Reber (1989) which suggested set-wise deep structure is best described as any set of rules that can generate the grammatical items. Reber (1989) suggested this because he believed that items generated from a grammar carried the general structure of the domain, but did not exhibit any important structure in and of themselves. However, it can be argued that items do have two very important structural aspects, and that the nature of set-wise structural regularities can be described as the frequency with which specific aspects of item structure occur across items. Items have both surface structural and deep structural characteristics, and therefore, the set of grammatical items can be described as having both surface structural and deep structural regularities.

Surface structural information and deep structural information represent two general, variably independent forms of information made available by a set of

items generated from an artificial grammar. These two forms of information have different utilities with respect to their support of grammatical sensitivity under different types of tests. However, either of these two forms of information may or may not be incorporated into memory, in various possible forms of representation, after encountering a set of grammatical items. The sensitivity of human memory to these two forms of information will be discussed and empirically examined with respect to the episodic-processing account of implicit learning.

The Episodic-Processing Account and Sensitivity to Structural Regularities

The episodic-processing account predicts a number of ways in which people can become sensitive to the set-wise structural regularities of grammatical items as measured by same and changed letter-set classification tests. These predictions rely upon the differential utility of different forms of deep structural and surface structural information outlined above. Equally important, these predictions rely upon the different processes memory can perform on the different aspects of the structure of individual grammatical items. All of these predicted forms of processing are hypothesized to support grammatical sensitivity under different circumstances, without invoking an automatic abstraction routine which extracts information across exemplars and stores them in a summary representation. All of these predicted forms of processing require that memory process different aspects of stimulus structure under the guidance of an incidental purpose given by the experimenter, form a representation of these processing experiences, and retrieve these experiences to support later performance under similar circumstances.

First, the sensitivity of memory to the set-wise surface structural regularities of a set of grammatical items will be discussed. This will be followed by a discussion of memory's potential sensitivity to set-wise deep structural regularities.

Sensitivity to Set-Wise Surface Structural Regularities. The first way in which the episodic-processing account suggests that memory can become sensitive to the set-wise structural regularities of a set of grammatical items is predicted to transfer only to items based upon the same letter-set.

The sensitivity to the set-wise surface structural regularities of a set of grammatical items could be governed by representations of the surface structural sequences of individual grammatical training items. According to the episodic-processing view, by processing the sequences of surface structural features of individual items, individual representations of these processing-experiences are created. These representations may include surface structural sequences of variable lengths, ranging from a single surface structural feature, to bigrams, to the entire surface structural sequence of an item. However, these representations will be encoded in terms of the processing that was performed on these sequences. Therefore, according to the episodic-processing account, the similarity between current and past experiences will not be a function of number of matches between the structure aspects of the training and test items, predicted by the instance account (Vokey & Brooks, 1992) and the bigram knowledge account (Perruchet & Pacteau, 1990). Instead it will be guided by the psychological similarity between the representations of surface structural aspects of training and test items as they are processed for a particular purpose.

Experiments 1 and 2 were conducted to determine whether grammatical sensitivity is mediated by the objective structural similarity between training and test items or by the psychological similarity between training and test experiences.

Sensitivity to Set-Wise Deep Structural Regularities. The episodic-processing account also predicts that incidentally attending to and processing the deep structural features, the pattern of element repetition, of individual items could support sensitivity to the general, set-wise, structural regularities of a set of grammatical items. This form of knowledge is predicted to support grammatical discrimination in both same and changed letter-set classification tests because grammatical items, based upon any letter set, share more deep structural features with each other than with nongrammatical items.

The deep structural features found in the set of grammatical training items could be represented in memory in one of two ways depending on the form of processing performed by the subject. The first involves the independent coding of the deep structural features of individual items. When the task requirements induce subjects to focus on and process the deep structural features of individual items separately, the episodic-processing account predicts that subjects will code these deep structural features independently. Subjects would, thus, achieve analytic, set-wise, deep structural knowledge in the form of representations of individual deep structural features. This possibility was tested in Experiment 5.

Under other circumstances, such as when the task focuses the subject on the whole item, the episodic-processing account predicts that subjects will process

the deep structural features of individual items in an integral fashion. This would result in distributed representations of compounds of deep structural features found to co-occur in the training items. The possibility that changed letter-set transfer is supported by representations of the complete set of deep structural features co-occurring in individual items was tested in Experiment 3.

Representations of surface structural sequences of individual training items are, as stated previously, predicted to support strong transfer to grammatical items based on the same letter set. However, for this form of representation to support transfer across letter-sets an abstract analogy mechanism would be necessary. That is, relational or abstract similarity between the surface structural sequences of items coded in training and the surface feature sequences of test items would have to be employed. This would require that whole or nearly whole training stimuli were coded and stored in memory in a form that is more or less one-for-one with the stimulus. This seems likely to occur only when subjects are given considerable time to code training items and no highly elaborative coding operations are performed. Experiment 4 was conducted to determine whether, under conditions in which the complete coding of whole items is improbable, surface structural coding would support grammatical transfer under changed letter-set conditions. The episodic-processing account predicts that this form of coding will not support changed-letter set transfer, although it will support same letter-set transfer.

To reiterate, conditions under which subjects are incidentally led to code the surface structural sequences of individual training items, should, according to the episodic-processing account, result in good transfer to items based on the

same letter-set and poor transfer to items based on a different letter-set. Such coding conditions serve as a strong test between the episodic-processing account and Reber's (1989) structural abstraction account. Reber's account proposed that memory automatically abstracts deep structural regularities, in the form of set-wise rules, when subjects are not encouraged to code instances. Reber also argued that changed letter-set transfer is the best test of this abstraction. Therefore, it is clear that Reber's structural abstraction account, in opposition to the episodic-processing account, predicts that both same and changed letter-set transfer should always occur under situations when the subjects are not encouraged to learn instances. Experiment 4, in which the incidental demands of the training task did not encourage subjects to memorize training items, had all the requirements necessary to serve as a test between these two accounts of implicit learning.

Summary. As a whole, the current experiments were designed to demonstrate the power of the episodic-processing account to explain the variability of performance that occurs when subjects experience exemplars of an artificial grammar for different incidental purposes. Different incidental purposes for encountering exemplars of a grammar should cause critically different operations to be performed on particular aspects of stimulus structure. To foreshadow the results, manipulation of the demands of the task in which exemplars were incidentally encountered radically modified the nature of the knowledge base acquired in implicit learning. It was concluded that accounts suggesting that grammatical sensitivity is a result of automatic abstraction of set-wise structural regularities, or that any specific form of stimulus structure

is primary, are insufficient to explain the variability of performance found in implicit learning.

Experiment 1

The purpose of Experiment 1 was to test whether memory is automatically sensitive to structural regularities, or whether structural sensitivity is mediated by specific representations of particular processing experiences. Do abstraction routines automatically extract the underlying invariance patterns of the stimulus environment? Or, does the processing performed on stimulus structure form part of the representations which support grammatical sensitivity? Most of the accounts outlined in the introduction assume that structure is automatically acquired by memory. Reber (1989) was especially dubious of any account of implicit learning which suggested that the elaborative operations performed by the subject have anything but a superficial impact on grammatical sensitivity.

To achieve the purpose of this experiment subjects had to be given different task to perform in dealing with the exemplars of the grammar. Two different form of incidental repetition were chosen. The first, string repetition, involved repeating the letters of an item from left to right. For example, the item MTC would be repeated out loud "MTC, MTC, MTC, MTC". The second, letter repetition, involved repeating each letter of the item four times. For example, the item MTC would be repeated out loud "MMMM, TTTT, CCCC". These forms of repetition may or may not affect the representation of grammatical knowledge supporting grammatical sensitivity.

To ensure that the processing was totally incidental with respect to the fact that there was a grammar and incidental to the fact that the training items were of any importance, the incidental learning procedure developed by Rundus (1977) and Glenberg, Adams and Smith (1977) was used. This procedure involved informing subjects that they were in a number learning experiment and that repetition of the letter strings was "a distraction". This procedure ensured that repeating the letters of the items did not involve any special elaborative processing. Reber (1989; Reber & Allen, 1978) argued that any undue elaborative processing of the training instances causes interference with default abstractive processing. In fact, this form of incidental grammar learning, in the absence of knowledge that the grammatical training items are of any importance whatsoever, can be considered less intrusive than even the observation condition used by Reber and Allen (1978), which cued subjects that items were of primary importance.

The critical aspect of this experiment involved determining whether reinstating the form of processing performed on the training items affected grammatical sensitivity at test. Subjects either performed the same type of repetition on all training and test items (matched conditions), or performed one type of repetition on the training items and the alternative type of repetition on the test items (unmatched conditions). If grammatical training items are represented in terms of the processing performed by memory, then classification should be selectively influenced by which form of processing was required at test. The episodic-processing account suggests that classification should be assisted when the functional representation of current experience is highly

similar to past experiences. Functional representations include the particular aspects of stimulus structure that have been processed, in the form they have been processed. In the matched processing conditions, the functional representations of grammatical items processed in training should achieve maximal similarity to the functional representations of the grammatical test items. In this case, not only would the structure of the grammatical training and test items be similar, but their representations, including the way in which they were processed, should be similar. In contrast, past experience should be less successful in guiding classification in the unmatched conditions because the functional representations of the test items, whether those items are grammatical or nongrammatical, should be highly dissimilar to the functional representations of grammatical training items because the form of processing applied at training and test are dissimilar. As a result, subjects who perform matching repetition operations at training and test should receive the maximum benefit from experience.

Therefore, the episodic-processing account predicted that in the matched conditions, when the form of repetition performed remained the same between training and test, classification performance should be better than in the unmatched conditions, when the form of repetition performed changed between training and test. Structural accounts of grammar learning, whether they are bigram accounts (Perruchet & Pacteau, 1990), instance accounts (Vokey & Brooks, 1992), or structural abstraction accounts (Mathews et al., 1989; Reber, 1989) would predict that there should be no such effect of matching processing contexts because only the structure of the training and test items

should influence performance. That is, if subjects acquire purely structural information when exposed to grammatical items, then subjects should be able to discriminate between grammatical and nongrammatical test items regardless of the way they were processed.

Method

Subjects. Forty undergraduates participated for course credit.

Materials. The grammar shown in Figure 3 was used to generate all grammatical items. This grammar and all the grammatical and nongrammatical items used in this experiment were created and used by Reber and Allen (1978). These items were also used in experiments conducted by Dienes et al. (1991), Dulany et al. (1984) and Perruchet and Pacteau (1990). It should be noted that one change was made to the items. The letter V was consistently replaced with the letter C to remove the influence of the highly salient bigram TV and trigram MTV.
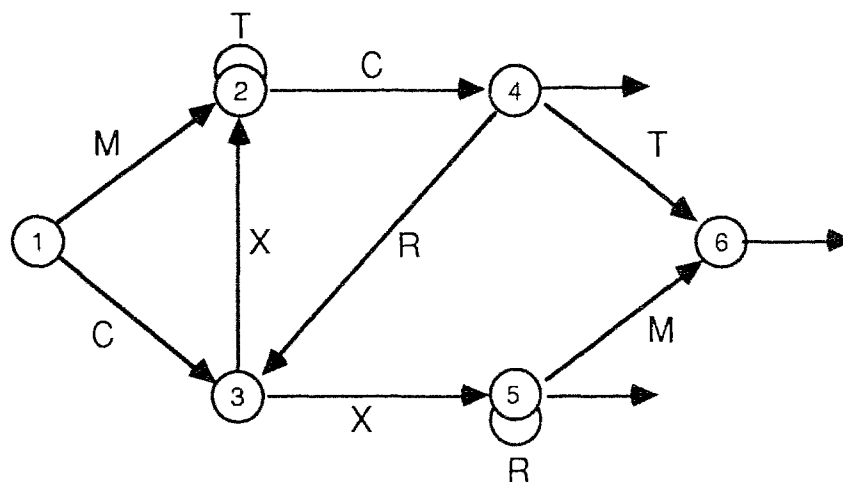
Figure 3. The finite state grammar used in Experiment 1 and 2.

A total of 20 items were presented in training. A total of 25 grammatical items were presented at test: 5 which had been presented previously in training, and 20 which had not. Also presented at test were twenty nongrammatical items created by introducing at least one grammatically impermissible letter substitution into a grammatical item. (See Appendix A for a complete listing of all training and test items.)

The experiment was conducted on a Macintosh IIci computer. All stimuli were presented in the center of the monitor using a Monaco, 24 point font.

Procedure. The procedure used in this experiment was designed to determine whether reinstatement of distinctive processing performed on letter strings in training had any affect upon classification performance at test. To this end, two distinctive methods of repeating the letter sequences were introduced. The first, string repetition, required subjects to repeat each letter of the item out loud from left to right. They had to do this four times. For example, subjects performing string repetition on the item MTRM would repeat out loud "MTRM, MTRM, MTRM, MTRM". The second, letter repetition, required subjects to repeat each letter of the item out loud four times in succession. They proceeded from left to right. For example, subjects performing letter repetition on the item MTRM would repeat out loud "MMMM, TTTT, RRRR, MMMM". Subjects were required to perform one of these types of repetition at training and one of these types of repetition at test. This resulted in a 2 X 2 between-subjects design, with string or letter repetition at training crossed with string or letter repetition at test. Therefore, there were four between-subjects conditions: string/string repetition; string/letter repetition; letter/string repetition,

and; letter/letter repetition. Ten subjects were randomly assigned to each condition.

The training procedure was designed to induce subjects to encode the letter sequences of the training items in a distinctive manner without being aware that they were doing so, or being aware that the items were constructed according to a set of rules. To accomplish this, following the procedure designed by Rundus (1977) and Glenberg, Smith and Green (1977), subjects were informed that they were in a number learning experiment measuring short-term memory and that they were to repeat out loud a row of letters as a distraction task. The "row of letters" was, in fact, a grammatical training item. On each training trial a three digit number appeared on the screen for three seconds. Subjects believed remembering this number was their main task. The number was then replaced by a grammatical training item and subjects were required to repeat it out loud in the manner appropriate for the condition they were assigned to. Upon completing the repetition, the subject hit the space bar on the keyboard and a prompt appeared on the screen asking for the three digit number. After the subject entered the number, there was a one second delay and then the next trial began. The whole set of 20 grammatical training items were presented three times in three different random sequences that were re-randomized for each subject.

Immediately following the training session, subjects were informed that all of the letter strings they had repeated were generated by a complex set of rules which allowed only certain letters to follow other letters. They were informed that they would now see more letter strings and that they would have to judge

whether they followed the same rules for letter ordering as the strings they had originally repeated. In addition, they were told half of the letter strings they would judge followed the rules and half did not. In accordance with the condition they had been randomly assigned to, subjects were also instructed how to perform the type of repetition they had to perform at test.

At test, each of the 25 grammatical items and 25 nongrammatical items were presented in a random sequence and then again in a different random sequence. This test procedure, involving the use of two test phases incorporating the same set of test items, has been used previously with these same stimuli by other researchers (Dienes et al., 1991; Dulany et al., 1984; Perruchet & Pacteau, 1990; Reber and Allen, 1978). Subjects had to repeat the item in the way required and then classify it, as either grammatical or not, by entering <Y> for "Yes, it follows the rules" or <N> for "No it does not follow the rules". After entering their response, the next item was presented after a one second delay. At the end of the test session, subjects were asked: 1) whether they had suspected during training that they were going to be tested later on the letter strings; 2) how they had attempted to perform the test task. They were then fully debriefed on the nature of the study.

Results and Discussion

The episodic-processing account predicted that experiences would support current processing to the extent that the representations of the current and prior stimuli, as processed, were similar (Whittlesea, 1987; Whittlesea & Brooks, 1988; Whittlesea & Cantwell, 1987). In the matched repetition cases, not only the stimulus structures of grammatical items were similar, but the way in

which that structure was being processed was similar. This was predicted to result in maximum grammatical sensitivity.

For each training stimulus, the mean time taken to perform string repetition ($\underline{M}$ = 8.0 seconds) and letter repetition ($\underline{M}$ = 7.7 seconds) were not reliably different, t(38) = .708, p = .443. This means that any differences in the two conditions cannot be a result of the duration of the training session.

Table 1
Experiment 1: Mean Proportion of Correct Classification Responses as a Function of Experimental Condition and Test Phase

Proportion Correct

Group	First Half	Second Half	Total
String/String	.68	.61	.65
String/Letter	.61	.56	.59
Letter/String	.56	.52	.54
Letter/Letter	.59	.61	.60

Taking the proportion of correct responses for the entire set of test items as the dependent measure, an analysis of variance (ANOVA) was performed to determine the effects of the type of training repetition and the type of test repetition. As can be seen in Table 1, subjects who performed string repetition in training were better at classifying test items than were subjects who performed letter repetition at training, with this main effect of training being reliable, $\underline{F}$(1,36) = 9.881, $\underline{p}$ = .003. From an episodic-processing standpoint,

this effect is likely the result of subjects in the string repetition condition coding longer sequences of the training items than did subjects in the letter repetition condition. In the string repetition condition, the entire string was processed in a serial left-to-right fashion because the entire string was the unit to be repeated four times. However, in the letter repetition condition, subjects coded each letter of the training string in a relatively separable fashion because a single letter was the unit to be repeated four times, before continuing. The relatively independent coding of individual letters in the letter repetition condition would lead to less sensitivity to set-wise surface structural regularities because the frequency with which individual letters occur in this set of items was not as predictive of grammatical status as were longer sequences (cf. Perruchet & Pacteau, 1990). In terms of an automatic structural abstraction account, this string repetition advantage could only be explained by arguing that letter repetition is an unusual form of processing which disrupts the default tendency to unconsciously abstract general structure.

There was no reliable main effect of the form of repetition performed at test, $F(1,36) = .001$, $p = .978$. However, there was a reliable interaction between the effects of training and test repetition, $F(1,36) = 10.958$, $p = .002$, with performance in the matched repetition conditions being higher than in the unmatched conditions. Breaking down this interaction, the proportion of correct test responses for the subjects who performed string repetition at training and test (string/string condition; $M = .65$) was reliably higher than for the subjects who performed string repetition at training and letter repetition at test (string/letter condition; $M = .59$; $t(18) = 2.209$, $p = .040$). Moreover, the

proportion of correct test responses was reliably higher in the letter/letter condition ($\underline{M}$ = .60) than in the letter/string condition ($\underline{M}$ = .54; $\underline{t}$(18) = 2.506, $\underline{p}$ = .022). As suggested by the episodic-processing account, these effects of the reinstatement of the processing context at training and test appeared to support the conclusion that the knowledge base supporting grammatical discrimination preserves the specific form of processing performed on the training items.

As can be seen in Table 1, performance was not stable across the two passes through the test stimuli. In fact, in every condition, except the letter/letter condition, classification performance was less accurate in the second test. An analysis of variance (ANOVA) including test phase (first or second) as a factor showed that this general deterioration in performance was marginally reliable, $\underline{F}$(1,36) = 3.629, $\underline{p}$ = .065. Although others have used the same stimuli and testing procedure, a drop in classification accuracy in a second pass through test stimuli has not been previously reported (Dienes et al., 1991; Dulany et al., 1984; Perruchet & Pacteau, 1990). However, these researchers all used standard memorization training. Therefore, this effect is likely a result of the training procedure used in the current experiment. The incidental nature of coding of the training items could have resulted in only very inefficient encoding. Such inefficient encoding would be open to interference from test experiences. That is, repeating the test items and judging them in the first test phase would create further codes of individual items. These codes, of both grammatical and nongrammatical test items, might then interfere with performance on the second test by being retrieved, along with the codes of training experiences, and guiding

performance. However, they would be of both grammatical and nongrammatical training items, and therefore, would likely decrease grammatical sensitivity by making both grammatical and nongrammatical items available. This argument is supported by an experiment by Dienes et al. (1991), which demonstrated that subjects who memorized both grammatical and nongrammatical training items performed more poorly in a later classification than subjects who memorized only grammatical training items. It would appear that the coding of nongrammatical items is capable of interfering with grammatical sensitivity. This same argument can also explain why the deterioration in performance in the second test phase was not discovered in previous studies using standard memorization instructions. Under standard memorization conditions this same process could occur, but the richer, more elaborate codes of training experiences would likely minimize the effect of inefficiently encoded test experiences.

Because of the deterioration in performance in the second test phase, the data were reanalyzed using only the data from the first pass through the test stimuli; these data could not be influenced by the repetition of test items. For the first-half data, the subjects who performed string repetition in training were still reliably better at classifying test items than were subjects who performed letter repetition at training, $F(1,36) = 8.740$, $p = .005$. Again, there was no main effect of the form of repetition performed at test, $F(1,36) = .335$, $p = .558$, but there was a reliable interaction between the effects of training and test repetition, $F(1,36) = 10.958$, $p = .002$. Breaking down this interaction, the proportion of correct test responses in the string/string condition ($M = .68$) was higher than in the string/letter condition ($M = .61$),

but this effect was only marginally reliable, $t(18)= 1.972$, $p = .064$). In addition, the proportion of correct test responses was not reliably higher in the letter/letter condition ($\underline{M} = .59$) than in the letter/string condition ($\underline{M} = .56$; $t(18) = 1.134$, $p = .272$).

Therefore, subjects who performed string repetition on training did get some benefit from performing string repetition again at test, but there was no similar effect for letter repetition. The latter result can be explained by both the episodic-processing account and the structural abstraction account. From the episodic-processing perspective, it is highly believable that subjects who performed letter repetition at training did not receive much support from reinstatement of this processing at test because the form of processing they performed did not result in a great deal of grammatical discrimination under either test condition. Dulany et al. (1984) had a control group perform this same classification test without training and they achieved 55% accuracy. Therefore, the amount of learning demonstrated by subjects performing letter repetition at training (mean percent correct = 58%) was not of great magnitude. As stated previously, it is likely subjects performing letter repetition at training coded the individual letters of the training items in a relatively independent fashion. Therefore, requiring subjects to perform this same form of coding at test would not create a great deal of benefit, simply because it is inefficient.

A structural abstraction account could explain these first-pass results for letter repetition by arguing that this form of repetition disrupts both the acquisition and utilization of grammatical knowledge. That is, the fact that

subjects in the string/string condition performed better than the subjects in other three conditions could be a result of the disruptive effects of letter repetition in those other conditions. This could be the case because is it a less practiced form of rehearsal. However, a structural abstraction account would have difficulty in explaining why the accuracy of classification performance dropped in the second pass through the stimuli. Structural abstraction of set-wise regularities has been assumed by Reber (1989; Reber & Allen, 1978) to result in stable representations capable of supporting consistent back-to-back responses to the same set of grammatical test items. The reason why this would not be the case, especially in the string/string condition, is not obvious.

Experiment 2

The episodic-processing account has the ability to explain a little-cited result in the artificial grammar learning literature. Gordon and Holyoak (1983) reported that after standard training, subjects, who were unaware that there was a grammar, rated grammatical test items as more pleasant than nongrammatical test items. This extension of the mere exposure effect (vide Zajonc, 1980) to the issue of grammatical sensitivity is not predicted by most theories of implicit learning. In fact, it links two seemingly unrelated findings on implicit memory performance.

The episodic-processing account, however, suggests that this should be the case, and as well, that recognition (as demonstrated by Vokey & Brooks, 1992) could serve as a task in which grammatical sensitivity can be demonstrated. The episodic-processing account bases its explanation on the concept of processing

fluency (Whittlesea, 1992), which incorporates the notion of perceptual fluency as proposed by Jacoby (Jacoby & Dallas, 1981; Jacoby & Whitehouse, 1989; Whittlesea et al., 1990). It is hypothesized that the ease of processing a stimulus creates a feeling of familiarity which can serve as the basis for recognition decisions. This ease of processing can be caused by the retrieval of past experiences or some current source not known to the subjects, such as the density of a noise mask which influences the ease of perceiving a target word (Whittlesea et al., 1990). The nature of the feeling state created by the fluency in processing is hypothesized to be governed by an unconscious attribution process (Jacoby & Whitehouse, 1989; Whittlesea, 1992; Whittlesea et al., 1990).

With respect to the mere exposure effect, Whittlesea (1992) demonstrated that both the prior presentation of a word and also a covert manipulation of processing fluency increased the subjective likability of a target word. In that experiment, the processing fluency of a target word was increased, as measured by a decrease in pronunciation latency, by a past experience with the tar: t or, alternatively, by the target's current semantic predictability in a phrase. If the processing fluency of a target word was increased, the subjects were more likely to claim that it was old or more likeable.

Attributions of processing fluency could also support decisions of grammaticality. The past processing of a similar item in a similar way could increase the fluency of processing a grammatical test item making it, depending upon the context and the decision to be made, seem alternatively, grammatical, likeable, or old. Prior experience of processing similar grammatical items

would support the processing of grammatical test items, thereby increasing the fluency of that processing. Nongrammatical test items would not receive as much support from past experience, and therefore would not be processed as fluently as grammatical items. Therefore, the fluency of processing the test items would result in feeling states correlated with the grammatical status of the test items.

Therefore, the episodic-processing account predicts that the same mechanism underlies decisions made by subjects in the seemingly divergent cognitive paradigms of recognition, implicit learning and mere exposure. These decisions have been theorized to require very different cognitive mechanisms, such as an intra-item integration process for recognition (Mandler, 1980), a separate affect system for mere exposure (Zajonc, 1980), and an unconscious structural abstraction system for grammatical sensitivity (Reber, 1989). However, the implicit processes underlying these divergent decisions are, according to the episodic-processing account, highly similar. Therefore, Gordon and Holyoak's (1983) finding that grammaticality of test stimuli affects their likability is not surprising from the episodic-processing standpoint.

Because of the close relationship between the fluency of processing and recognition decisions (Whittlesea, 1992; Whittlesea et al., 1990), and the possibly close relationship between recognition decisions and classification decisions (e.g., Brooks & Higham, 1992; Whittlesea, 1983), a recognition decision was used as an indirect measure of grammatical sensitivity in Experiment 2. A recognition decision was used in an attempt to create a condition in which grammatical knowledge could be applied more implicitly. Subjects do not have to be informed that there is a complex set of rules underlying the

stimuli when a recognition decision is used. This change could reduce the use of explicit rule-generating strategies which might be encouraged by telling subjects that there is a complex set of rules, and that they must determine if test items follow those rules or not. As a result, subjects might be induced to rely on nonanalytic fluency-based information to a greater degree when making recognition decisions than when making classification decisions. This should serve to increase the effect of matching the processing contexts between training and test because this effect is hypothesized to act implicitly.

The procedure for Experiment 2 was identical to the procedure of Experiment 1, except that a recognition decision was used in place of the more usual classification decision. Subjects attempted to recognize a set of items that, except for five grammatical items, were all novel. To score the data, false alarms in the form of claimed recognition of novel grammatical and nongrammatical items was the primary measure. It was predicted that proportion of false alarms for grammatical items would be greater than the proportion of false alarms for nongrammatical items. Moreover, this difference was predicted to be greater in the matched processing conditions than the unmatched. The matched processing between training and test should result in grammatical items being more similar to the experiences of grammatical training items, which should increase the fluency of processing of those items, and in turn, result in them feeling more familiar. However, the nongrammatical items, which were not as structurally similar to grammatical training items, should get less support from past experience. In contrast, in the unmatched condition, both grammatical and

nongrammatical items should seem unfamiliar because they are both dissimilar to past experience in terms of the processing performed.

Method

Subjects. Thirty-two undergraduates participated for course credit.

Procedure. Unlike the instructions given to subjects before the test in Experiment 1, subjects were not informed that the letter strings that they had repeated out loud earlier conformed to a complex set of rules. Instead, subjects, before test, were simply informed that they were in a study which now required them to recognize the letter strings that they had repeated out loud during the number recall session. At test, subjects responded <O> for "Old" and <N> for "New". In all other ways, the procedure and materials used in this experiment were identical to that of Experiment 1. Eight subjects were randomly assigned to each condition.

Results and Discussion

In this experiment, sensitivity to the structure of the grammar would be indicated if grammatical items were classified as old more often than nongrammatical items. To allow direct comparison with the results of Experiment 1, which were reported in terms of the proportion of correct classification responses, the total number of the grammatical items classified as old was summed with the total number of nongrammatical items classified as new and then divided by the total number of responses. This indirect measure of grammatical sensitivity was calculated for each subject. As can be seen in Table 2, subjects judged grammatical items to be old and nongrammatical items to be new more often than would be expected by chance. Recognition decisions can,

therefore, serve as an indirect measure of grammatical sensitivity because items

conforming to the structure of the grammar seem more familiar than those that

do not.

In this experiment there were actually 5 old grammatical items among the 50

items presented at test. These items were included to allow comparison of these

results to the results of Experiment 1 and studies performed by other authors.

However, these items did not affect the results presented. The claimed

recognition of old grammatical items did not differ considerably from claimed

recognition of novel grammatical items in each condition.

Table 2
Experiment 2: Sum of the Grammatical Items Classified as Old and
Nongrammatical Items Classified as New as a Proportion of Total Test
Responses

	Test Responses		
Group	First Half	Second Half	Total
String/String	.70	.61	.66
String/Letter	.62	.63	.62
Letter/String	.59	.50	.55
Letter/Letter	.61	.52	.57

However, performance was not stable across the two passes through the test

stimuli. In fact, in every condition, except the string/letter condition, the

indirect measurement of grammatical sensitivity was lower in the second test.

An analysis of variance (ANOVA) including test phase (first or second) as a factor showed that this general deterioration in grammatical sensitivity was reliable, F (1,28) = 14.348, p = .001. Because of this effect, discussed in Experiment 1, the data from the subjects second pass through the test stimuli was not included in the analysis. The data collected during the subject's first pass through the test stimuli could not be influenced by the repetition of test items.

For the first pass through the stimuli, the subjects who performed string repetition in training were more sensitive to the grammatical status of test items than were subjects who performed letter repetition at training, $F(1,28)$ = 4.719, p = .038. There was no independent effect of the form of repetition performed at test, $F(1,28)$ = 1.335, p = .258. However, there was a marginally reliable interaction between the effects of training and test repetition, $F(1,36)$ = 3.950, p = .057. Breaking down this interaction, the indirect measure of grammatical sensitivity was reliably greater in the string/string condition (M = .70) than in the string/letter condition (M = .62), $t(14)$= 2.427, p = .029. However, the indirect measure of grammatical sensitivity was not reliably greater in the letter/letter condition (M = .61) than in the letter/string condition (M = .59, $t(14)$ = .546, p = .594). These results directly mirror the results presented in Experiment 1 for classification judgments.

The results of this experiment are neutral with respect to deciding between the episodic-processing account and the structural abstraction account. Both accounts can explain these results in the same way they could explain the results of Experiment 1. A structural abstraction account could explain these results by

arguing that letter repetition disrupts both the acquisition and utilization of grammatical knowledge. From the episodic-processing perspective, letter repetition appeared to be an inefficient method of processing the training stimuli. Letter repetition likely caused subjects to independently code the individual letters of the training stimuli. This relatively independent coding of individual letters in the letter repetition conditions would lead to only a small amount of sensitivity to set-wise surface structural regularities because the frequency with which individual letters occur in this set of items was not highly predictive of grammatical status. Moreover, requiring subjects to perform this same form of coding at test would not create a great deal of benefit, simply because it is inefficient. Therefore, it appears that the inclusion of letter repetition nullified the chances of Experiments 1 and 2 being strong tests between the episodic-processing account and the structural abstraction account.

This argument, suggesting that letter repetition was a poor form processing to use in the current experiments, is supported by a study by Whittlesea and Dorken (1992, Experiment 3) which demonstrated that grammatical sensitivity does appear to be supported by highly specific representations of training experiences. In a set of studies, with the same purpose as the current Experiment 1 and 2, subjects performed one of two types of processing on each training and test item. These two tasks both involved two different methods of complete, integrative coding of the surface structural sequences of the training items. In this way, they are better processing tasks than the letter repetition used in the current experiments, which involved relatively separable coding of each letter of a stimulus. Specifically, a spelling task, similar to string

repetition used in the current experiments, was used in conjunction with a pronunciation task. Using these two tasks, it was demonstrated that classification performance was selectively influenced by the processing requirements demanded at test. When subjects processed a test item in the same manner they had processed the training items of a grammar, they were more sensitive to its grammatical status. It was concluded that the similarity between current and prior experiences, rather than the similarity between current and prior stimulus structures, guided performance. Therefore, the inconclusive results of Experiment 1 and 2 were probably due to the relatively separable nature of coding operations performed on individual letters by subjects in the letter repetition conditions.

Experiment 3

Observations of transfer across changes in the letter-set instantiating the items of the grammar have been argued to support accounts of implicit learning postulating that memory forms representations of the underlying rules of the grammar (Mathews et al., 1989; Reber, 1969, 1989). According to these accounts, grammatical knowledge is summarized in memory by abstract representations that are formed by extracting set-wise information from the instances presented. Reber (1989) suggested that these abstract representations directly represent the general rule-based structure of the domain. In contrast, Brooks and Vokey (1991) have argued that a memory base of stored instances could serve changed letter-set classification if an abstract similarity mechanism was used at test.

With respect to changed letter-set transfer, the episodic-processing account predicts a number of ways that subjects can become sensitive to deep structural regularities. Two of these involve relatively complete knowledge of individual training items. The first is the abstract analogy mechanism put forward by Brooks and Vokey (1991). When subjects are directly or indirectly induced to code the entire surface structural sequences of individual training items, they should form representations of those surface structural sequences. Then, at test, the abstract or relational similarity between the current and past sequences should guide successful classification performance. The second way that relatively complete knowledge of individual training items could support sensitivity to deep structural regularities involves subjects coding the complete set of deep structural features, the specific repetition patterns of individual elements, within individual training items. The later use of the matches between these specific deep structural patterns, as encoded, and the deep structural patterns of individual test items could support generalization to novel grammatical items instantiated upon a different letter set. This is possible because grammatical items, based upon any letter set, share more sets of deep structural features with each other than with nongrammatical items. The coding of an item's complete set of deep structural features should result when the induction task directly or indirectly induces subjects to code the complete pattern of repetition within individual training items.

Representations of the deep structural features of individual training items should be more efficient than representations of surface structural sequences in supporting transfer to changed letter-set stimuli. Deep structural pattern

information is directly relevant for success in a changed letter-set test, but surface structural information is only indirectly relevant through the use of an abstract analogy mechanism. Both forms of coding require the later use of similarity, but the abstract analogy mechanism involves postulating memory's additional ability to compute similarity on an abstract or relational dimension. This indirect computation of similarity would require an extra step in the calculation of similarity. Specifically, a translation process would have to match the past and current experiences on an abstract dimension before similarity could be computed.

Therefore, the episodic-processing account suggests that there are two ways that subjects could become sensitive to the deep structural patterns of individual items. The episodic-processing account is directly opposed to the structural abstraction account's suggestion that grammatical sensitivity stems from general knowledge of the deep, abstract properties of the grammar induced across items. Critically, such abstraction accounts, by their very nature, could not explain why subjects are sensitive to the deep structural patterns of individual items, if such sensitivity could not be supported by set-wise, abstract knowledge. However, the episodic-processing account explains such sensitivity as the result of the mechanisms outlined above. Therefore, to test between these two types of views, a changed letter-set transfer test, the traditional test of the abstractness of knowledge (Mathews, 1991; Reber, 1969, 1989), was used to determine whether subjects were sensitive to the deep structural patterns of individual items, when such sensitivity could not be supported by set-wise, abstract knowledge.

To achieve this purpose, a "grammarless grammar" was constructed. All stimuli in this experiment were generated by a set of rules specifying that there must be a unique pattern of deep structural features within each individual item, but no overall pattern of deep structural features across the set of items. In training, subjects studied items that each carried an unique deep structural pattern. For this purpose, one half of the deep structural patterns generated from the "grammarless grammar" were instantiated upon the letter set {H,D,F,X,M,S,B,R}. The resulting items were presented to subjects given standard memorization instructions. At test, the whole set of deep structural patterns was presented, but instantiated upon the changed letter-set {C,P,K,T,Z,V,G,L}. Subjects were required to recognize the items they had studied in training, even though they were translated into a different letter-set.

The procedure for generating the individual deep pattern structures was such that no general information in the form of a prototype or a set of rules smaller than the number of patterns could assist the recognition of training patterns. Therefore, above-chance performance on the test task could not be supported by set-wise knowledge of abstract properties induced across experience with the set of training items. Subjects could only succeed by gaining fairly complete knowledge of individual training items and then utilizing this knowledge to recognize the same patterns at test. It is highly likely that complete item coding was encouraged by the standard request to memorize the items for a later memory test. Given such instructions, subjects likely attempt to learn as much as possible about each training item (cf. Reber, 1989, for an opposing view).

Method

Subjects. Seventeen undergraduates participated for course credit.

Materials. A single deep structural feature is defined as the specific

repetition pattern of a single element. For example, a single deep structural

feature of the item CXPF-XFCP can be expressed as "the first element is repeated

in the seventh position". The set of deep structural features, or the deep

structural pattern, of an individual item may be represented by the statement

that there must be a repetition of first element in the seventh position, the

second in the fifth, the third in the eighth, and the fourth in the sixth, or more

simply, with the arbitrary use of numbers as elements, as 1234-2413. A deep

structural pattern may be instantiated by substituting discrete elements of any

type, whether they are tones or lights of different durations or frequencies, or

letters of the alphabet.

For this experiment, thirty-two deep structural patterns were generated.

Each pattern contained four deep structural features. Specifically, four elements

were presented twice each in order to create each pattern (e.g., 1234-1423 and

1233-1442; See Appendix 2 for a complete list of stimuli). Each deep

structural pattern could share none, one or two deep structural features with any

other pattern. Each pattern, of course, shared four deep structural features

with itself, but could not share three deep structural features with any other

pattern.

The set of thirty-two patterns was divided into two sets of sixteen, one to be

used in training only, and one to be used in both training and test. The set of

sixteen training patterns were instantiated on letters randomly selected from the

letter-set {HDFXMSBR}. Thus, for example, the pattern 1234-1243 could be presented as HFXR-HFRX, or alternatively, as SMRX-SMXR. The translation between numbers and letters was entirely randomized between items and between subjects. At test, all items were randomly instantiated upon the letter-set {CPKTZVGL} using the same procedure. For example, the pattern 1234-1243 could be presented as CPKV-CPVK, or alternatively, as TZGK-TZKG. Therefore, the overlap in surface features was not a possible source of information.

The two sets of patterns were created by a Latin square procedure which ensured that the smallest set of rules that could specify either set of patterns was a set of sixteen rules specifying each pattern separately. This procedure also ensured that no pattern was more or less typical than any other pattern in its own set. Each pattern shared a total of eight deep structural features with the other members of its own set. Moreover, a high degree of consistency existed in the amount of overlap of deep structural features across the patterns in the training set. In fact, half of all patterns shared two deep structural features with two other patterns and one deep structural feature with four other patterns. The remaining half of the patterns shared two deep structural features with three other patterns and one deep structural feature with two other patterns. In addition, there was no clustering of similar patterns which would have allowed the set of patterns to be described by a compact set of rules. As a result, it was impossible to summarize the training set either through a compact set of rules or through a prototype, and these set-wise sources of information could not be used to discriminate repeated from nonrepeated patterns.

Another potential source of set-wise information was the analytic similarity between a test pattern and the set of training patterns. Analytic similarity is computed by counting the frequency with which the individual deep structural features of a test pattern occur in the entire set of training patterns. This source of information was not capable of discriminating between the two pattern sets, because the patterns were divided into sets to ensure that each deep structural feature occurred an equal number of times in both the repeated and nonrepeated test sets. This also ensured that any pattern in the repeated and nonrepeated test sets had identical analytic similarity to both sets.

The only remaining potential source of information for discriminating repeated from nonrepeated test patterns was nonanalytic similarity. Nonanalytic similarity is the overlap between a test pattern and a training pattern computed in terms of the compounds of deep structural features taken two, three or four at a time. However, compounds of deep structural features taken two at a time would be an insufficient source of information because non-repeated patterns shared two deep structural features with one of the repeated patterns as often as a repeated pattern shared two deep structural features with one of the repeated patterns. Moreover, items could not match other items on three deep structural features. Therefore, only knowledge of the conjoint occurrence of all four deep structural features defining a pattern could discriminate between repeated and nonrepeated patterns.

Any three deep structural features, taken together, exclusively specified a single pattern. That is, once three deep structural features were specified there were no longer any degrees of freedom left for the construction of the pattern and

the final deep structural feature, by default, was also specified. Therefore, subjects did not need complete knowledge of the set of all four deep structural features of an individual pattern to recognize it as repeated, they only needed three. Thus only fairly complete knowledge about the set of deep structural features defining the construction of particular item could support discrimination between a repeated a non-repeated pattern.

Procedure. The experiment was conducted on a Macintosh IIci computer. All stimuli were presented in the center of the monitor using Monaco, 24 point font.

At the start of the training phase, the subjects were given the following instructions:

This is a simple memory experiment. You will see items made of letters, all of which will be consonants. These items will be eight consonants in length. Each item will be presented on the screen for 10 seconds, and then replaced by another item. There is a total of 16 items. Your task is to learn and remember as much as possible about these items.

Instructions of this type have been used previously by Reber (1976), Reber et al. (1980), Dulany et al. (1984), Perruchet and Pacteau (1990) and Dienes et al. (1991) in artificial grammar learning experiments.

Following the training instructions, each of the 16 training strings generated for that subject was presented on the screen for 10 seconds, with a one second interval between presentations. The entire set of training strings was presented a total of three times, in three different random orders each time.

The test phase immediately followed the training phase. The subjects were informed that the letter strings they had studied in training were generated by a complex set of rules. The subjects were told letter strings would appear on the screen, one at a time, and they would have to judge whether or not those strings

had been presented in training. The letter translation was carefully explained, using the strings AAEIEI and OOUYUY as a hypothetical example. When the subject understood the nature of the letter transformation, they were told to press "Y" on the keyboard if they thought "Yes" the item presented was one that had been presented in training, and "N" on the keyboard if "No" the item had not been presented in training. They were also informed that 16 of the 32 items had been presented in training.

The set of 16 patterns presented in training and the set of 16 patterns that had not been previously presented, all instantiated on the test letter-set, were used as test items. These items were presented on the screen, one at a time, in a random sequence. The prompt "Is this item Old? <Y/N>" remained at the bottom of the screen throughout the test phase. The subjects were self-paced, with each test item being presented immediately following the response to the last.

Following the test phase, the subjects were asked the following questions: 1) How did you attempt to learn about and memorize the training strings, and; 2) At test, given that the letters had been changed, how did you attempt to recognize the items you had seen previously? The subjects were encouraged to elaborate upon any thoughts they had concerning the two questions, and how their strategies may have affected their performance. The subjects were then fully debriefed about the nature of the study.

Results and Discussion

In principle, subjects could only discriminate old from new items on the basis of fairly complete knowledge of the conjoint occurrence of the deep structural features of individual items. They could not discriminate between old

and new items on the bases of any set-wise knowledge, such as rules or prototypes, because of the nature of the stimulus set used.

However, subjects were able to discriminate between repeated and non-repeated patterns with a mean accuracy of .57. This discrimination was reliably above chance, $t(16) = 3.490$, $p = .003$. There was no evidence that any single item or set of items was responsible for this above-chance discrimination. In fact, 10 of the 16 repeated patterns and 10 of 16 nonrepeated patterns were correctly classified by more than half of the subjects. Moreover, the pattern 1223-3441, which was presented last at test because of its salient repetition pattern, did not overly contribute to the above-chance discrimination. Mean accuracy (.56) was still reliably above chance when the data was reanalyzed without this item, $t(16) = 2.970$, $p = .009$.

Success also relied slightly more on hits ($p = .58$) than correct rejections ($p = .56$), consistent with the use of familiarity as a decision heuristic.

Subjects reported using many different strategies in an attempt to memorize the training items. For example, they reported attempting to: pronounce an item, by providing their own vowel sounds; repeat the letters of an item from left to right; compare the left and right sides of an item; look for the pattern of repetition of the individual letters; and make mnemonic phrases out of a whole item, or parts of an item. Most subjects reported using various strategies at different times during training because they were not certain that any one strategy was particularly effective. However, six subjects did report processing only the surface structural sequences of individual training items by repeating the letters from left to right or making words and mnemonic phrases from the

sequences of letters. Analyzed separately, these subjects performed at chance, achieving 50.0% accuracy. The remaining 11 subjects, who reported that they, at least once, attempted to analyze the deep structural features (i.e., code the repetition pattern of individual elements) of the items achieved 60.5% accuracy.

From the subjects' verbal reports it is apparent that subjects had many processing strategies available. Moreover, from the success of subjects who reported to, at least part of the time, code the deep structural features of individual items, it is apparent that some of these processing strategies were better than others with respect to performing well on a changed letter-set test. This correlation between the verbal reports of the subjects and their performance provides at least some tentative support for tne contention that coding the deep structural features of individual items is more useful in supporting later transfer to the changed letter-set test than is coding sequences of surface structural features and then attempting to use abstract analogies at test.

The most important conclusion from this experiment is that transfer across letter-sets does not necessarily indicate that the knowledge supporting this transfer is in the form of some abstract representation of general set-wise regularities, contrary to Reber (1969, 1989) and Mathews et al. (1989). However, in the current experiment changed letter-set transfer was supported by representations of the conjoint occurrence of deep structural features of individual items.

Moreover, in the more usual case, in which there are set-wise regularities underlying the set of grammatical stimuli, coding the deep structural features of

individual items could support grammatical sensitivity. Grammatical items necessarily share more deep structural features and conjunctions of deep structural features with themselves than with nongrammatical items. Thus, representations of the deep structural features of individual items could support grammatical sensitivity across changes in the letter set used to instantiate items. This does not mean, however, that it is not possible that summary representations of general set-wise regularities could not support sensitivity under those more usual circumstances. This is possible; however, it is not possible that such representations supported transfer in this experiment.

It should also be noted that the model THIYOS presented by Mathews and his associates (Druhan & Mathews, 1989; Mathews, 1990; Mathews, 1991; Mathews et al., 1989; Roussel et al., 1990) could not simulate the above-chance transfer demonstrated in this experiment. In THIYOS all condition-action rules were formed, strengthened and utilized in an independent fashion. Such condition-action rules were of individual features alone, and not the conjunctions of features underlying individual items. Therefore, THIYOS, which was substantially modified to simulate transfer across letter sets, would be insufficient to simulate these results.

Experiments 4-6

Experiment 3 demonstrated that the knowledge base which supports changed letter-set transfer does not necessarily have to represent set-wise structural regularities induced across instances. However, it did not test the contention made by Reber (1969, 1989) and Mathews et al. (1989) that, when an

underlying grammar is incidental to the demands of the induction task, memory automatically induces abstract representations of set-wise deep structural regularities. This was not tested because subjects could not have acquired set-wise knowledge in that experiment. Experiment 4-6 were designed and conducted to determine whether incidental exposure to grammatical training items, in a case where there is a grammar, automatically results in set-wise deep structural knowledge. Reber (1989) argued that changed letter-set transfer is the best test of this form of knowledge.

The episodic-processing account suggests that their are really two dimensions upon which set-wise structural regularities can form. As outlined in the general introduction, items have both surface structural and deep structural features. The frequency with which these features and groups of these features occur in the entire set of grammatical items describes the set-wise structural regularities. Set-wise, deep structural regularities can be described as the frequency with which each deep feature or conjunction of deep structural features occurs across the entire set of grammatical items. Set-wise, surface structural regularities can be described as the frequency with which certain surface structural features, or sequences of surface structural features, or sequences of surface structural features and their positions occur in the entire set of grammatical items.

According to the episodic-processing account, the form of knowledge that a subject acquires could be surface structural, deep structural, or both. Memory is capable of coding the surface structural and deep structural characteristics of grammatical training items; however, neither form of coding is automatically

employed. The knowledge acquired is dependent upon the specific processes that a subject performs on the training stimuli to achieve a particular purpose. The structural properties of the training stimuli will be experienced and represented in terms of the operations performed to accomplish a particular goal with respect to the stimuli. Therefore, subjects should become sensitive to surface structural regularities when they are directly or incidentally induced to code the surface structural properties of the grammatical training items. Similarly, they should become sensitive to deep structural regularities when they are induced to code the deep structural properties of the training items. In both cases, representations of particular aspects of item structure will support grammatical sensitivity when, at test, subjects process items with similar structural characteristics. This will support sensitivity to set-wise structural regularities because the frequency of occurrence of these structural features is more similar within the set of grammatical items than between the set of grammatical items and the set of nongrammatical items.

Experiment 4 was designed to test whether transfer to changed letter-set stimuli, the traditional test of deep structural sensitivity, is automatically supported by knowledge acquired while incidentally processing grammatical training items. To achieve this end, the incidental string repetition used in Experiments 1 and 2 was re-introduced as a method of processing the training items. Finding transfer to a changed letter-set classification test would support the structural abstraction account in its contention that deep structural regularities should be automatically abstracted whenever training items are incidentally encountered. However, the episodic processing account did not

predict changed letter-set transfer in this case. According to the episodic processing account, the demands of incidental string repetition should cause subjects to code the surface structural sequences of the training items. The resulting representations of surface structural sequences should not support changed letter-set transfer because items based on different letter-sets do not share surface structural features. However this form of surface structural coding should support same letter-set transfer because grammatical training and test items share surface structural sequences.

Experiment 5 was designed to demonstrate a case in which sensitivity to the deep structural regularities of a set of grammatical items could develop. A procedure was used to discourage processing of the surface structural sequences and encourage the processing of the deep structural features of the individual training items. To this end, subjects were required to judge whether each letter in the training items was repeated elsewhere in the item. These judgments were made as quickly as possible and the letter to be judged was indicated by a bar marker placed in a random order under each letter of an item, rather than from left to right. Moreover, the later test or the generating rules were not mentioned to the subjects; they were simply led to believe that the training task, designed to test their ability to rapidly identify letter repetition, was the only task in the experiment. Thus, subjects incidentally analyzed the deep structural features of the training items in the absence of any knowledge of the underlying grammatical rules or that the task had any relevance to a future task. It was expected that subjects would acquire deep structural knowledge of the training items. Specifically, it was expected that they would form representations of the

independent deep structural features occurring in the set of grammatical training items. It was predicted that this form of knowledge would transfer equally well to the same and changed letter-set transfer tests. This prediction was made because the deep structural features of the set of grammatical training items more closely matched their occurrence in the set of grammatical test items than in the set of nongrammatical test items. Any deep structural sensitivity in this case would result from the demands of the induction task, rather than the automatic abstraction of deep structural regularities across instances.

Experiment 6 was designed to allow comparison between the more usual memorization training procedure and the incidental training procedures used in Experiments 4 and 5. An intensive memorization procedure was used by Reber (1969) and Mathews et al. (1989) in tests of changed letter-set transfer. It is memorization training that is used most often in testing grammatical discrimination in tests of same letter-set transfer (e.g., Dienes et al., 1991; Dulany et al., 1984; Perruchet & Pacteau, 1990; Reber, 1976; Reber et al., 1980). In hopes of resolving the debate over what form of knowledge people acquire given this form of training, comparisons were made between the results of Experiment 6 and the other two experiments, in which the materials and tests were identical, but item coding was more tightly controlled at training.

Method

Subjects. Twenty-six undergraduate students participated for course credit in each of Experiments 4, 5 and 6.

Materials. The experiment was conducted on a Macintosh IIci computer. All stimuli were presented in the center of the monitor using a Monaco, 24 point font.

All grammatical stimulus strings were generated from the finite-state grammar shown in Figure 4. A total of fifty-five unique seven letter stimulus strings were generated from this grammar by starting at State 1 or State 2 and following the pathways of the grammar in the directions indicated until State 7 or State 8 was reached. To avoid large successive runs of the same element immediate reversals along a pathway were not allowed. For example, if the route between State 3 and State 6 was transversed, the reverse route from State 6 to State 3 was not allowed to immediately follow.


Figure 4. The finite state grammar used in Experiments 4-6 instantiated upon the training letter-set {M, T, C, X, R}.

Only 49 of the 55 strings generated from the grammar had unique deep structural patterns, that is, a unique set of deep structural features. An example a pair of strings sharing the same deep structural patterns are the strings RMRTXMX and MXMRTXT which both have repetitions in the first and third positions, the second and sixth positions, and the fifth and seventh positions. To create a set of strings all having unique deep structural patterns, 6 strings duplicating the deep structural patterns of other strings were discarded, leaving a set of 49 grammatical stimulus strings. Another 4 strings were arbitrarily discarded in order to create a set of 45 grammatical stimulus strings to be used in the experiment.

Twenty-five grammatical strings were selected to serve as training stimuli, and 20 grammatical strings were selected to serve as test stimuli. Care was taken in this selection process to ensure that the training items carried the set-wise structural regularities of the entire set of grammatical items, on both a surface and deep structural level.

Twenty nongrammatical stimulus strings were also created to be used at test. A nongrammatical string was defined as a string that did not have a deep structural pattern identical to any of the 49 unique particular deep structural patterns created by the finite-state grammar. This definition automatically ensured that the nongrammatical items did not share the same surface structural identity with any grammatical item. The nongrammatical strings were created by randomly changing one or two letters of a grammatical item used in training. This random replacement was governed by the following set-wise restrictions.

First, item by item each nongrammatical item matched a single grammatical test item in pattern complexity. For example, if a grammatical test item had one double repetition and one triple repetition (e.g., RMCXXMX), then a nongrammatical item was created with the same general pattern complexity (e.g., RMRTRTX). Second, both the grammatical and nongrammatical test items were created to have the same number of highly salient side-by-side repetitions of the same letter.

These two restrictions ensured that the set of test grammatical and set of test nongrammatical items had identical pattern complexity. If the nongrammatical items, item by item, had either more or less element repetition than grammatical items, it is conceivable that subjects could be successful at classification by simply using the strategy of categorizing items as grammatical if they have a great deal of repetition complexity. This strategy could not result in classification success with the current stimuli.

To instantiate each test item on a novel letter-set for the changed letter-set test, each of the original letters {M, T, C, R, X} used to instantiate the test items were replaced with a letter from the alternative {F, L, K, S, P} letter set. The mapping of letters was consistent within each item (e.g., all M's in an item would become F's), but was randomly reallocated between items (e.g., all M's in a different item could become L's). This letter translation was re-randomized for each subject. This form of randomization was used to eliminate the possibility that subjects could learn and utilize a consistent mapping between the old and new letter sets. (Cf. Mathews et al., 1989, regarding their claim that grammatical

sensitivity in the absence of knowledge of letter to letter mapping must be caused by "automatic abstraction of general knowledge of the grammar" (p. 1092).

Appendix B shows the complete set of training grammatical, test grammatical and test nongrammatical items instantiated upon the original letter set {M, T, C, R, X} and also an example of all grammatical and nongrammatical test items instantiated on the alternate letter set {F, L, K, S, P}.

Procedure. Three experiments were conducted, differing only in training conditions. Under all training conditions subjects were exposed twice to each of the 25 training strings. The whole set of training stimuli were presented on the screen, one at a time, in a random sequence, and then again, in a different random sequence. Therefore, there were a total of 50 training trials.

The training procedure for Experiment 4, the incidental repetition condition, was designed to require the subjects to repeat and encode the surface structural sequences of the individual training strings without knowing: 1) that the stimuli were generated by an artificial grammar; or 2) that there would be further testing of any type. To achieve this goal, an incidental learning procedure similar to the one developed by Rundus (1977) and Glenberg, Smith and Green (1977) was used. Subjects in the incidental repetition condition were informed that they were in a study measuring their short-term memory for three digit numbers. They were told they had to remember and then recall a three digit number presented on the screen. As a "distraction task", to make the number task more difficult, they were told they would have to repeat a series of letters out loud for 10 seconds before they recalled the number. In fact, this series of letters was a grammatical training stimulus. The sequence of events which

comprised a trial was: 1) a three digit number appeared on the screen for three seconds; 2) a blank screen for one second; 3) a grammatical training string was presented on the screen for ten seconds and the subject verbally repeat the letters from left to right as many times as possible; 4) the stimulus was replaced by a prompt to enter the three digit number, and; 5) the subject entered a three digit number. The next trial immediately followed. The subjects were led to believe that this task comprised the entire experiment.

The training procedure in Experiment 5, the incidental analysis condition, was designed to require the subjects to analyze and encode the deep structural features of the individual training strings without knowing: 1) that the stimuli were generated by an artificial grammar, or; 2) that there would be further testing of any type. In the incidental analysis condition, each training string was presented on the screen for 400 msec before a bar marker was placed under one of the letters. The subjects were required to hit the "A" key on the keyboard, with a finger on their left hand, if the letter underlined was repeated elsewhere in the stimulus string, and the "L" key, with a finger on their right hand, if the letter was not repeated elsewhere in the stimulus string. Subjects were told to respond with "one hundred percent accuracy, but also as fast as possible because both accuracy and response time are being measured." After a response, the bar marker disappeared for 200 msec and reappeared under a different letter in the stimulus, and remained on the screen until a new response was made. For each stimulus this procedure continued until the bar marker had been placed, in a random order, under each of the seven letters. The screen was then cleared and following a one second delay and the next trial began. Instructions to the

subjects at the beginning of this task simply involved telling them the decision and response they were to make upon seeing "an array of letters". The subjects were led to believe that this task comprised the entire experiment.

The training procedure for Experiment 6, the memorization condition, was similar to the training conditions used most often in artificial grammar learning experiments. The subjects in the Reber Memory Condition were given the following instructions:

This is a simple memory experiment. You will see items made of letters, all of which will be consonants. These items will be seven consonants in length. Each item will be presented on the screen for 10 seconds, and then replaced by another item. There is a total of 25 items. Your task is to learn and remember as much as possible about these items.

Instructions of this type have been used previously by Reber (1976), Reber et al. (1980), Dulany et al. (1984), Perruchet and Pacteau (1990) and Dienes et al. (1991). Each stimuli was presented on the screen for ten seconds, and then cleared from the screen for one second before the next stimuli was presented.

The test was identical for all three experiments. The test had two phases. In one, the same letter-set test phase, the letter-set used to instantiate the test items remained constant between training and test. In the second, the changed letter-set test phase, the letter-set used to instantiate the items was changed between training and test. The order of these test phases was counterbalanced between subjects. In each test phase, the subjects were exposed to the 20 grammatical and 20 nongrammatical test items instantiated upon the appropriate letter-set. The items were presented in a random order, held constant between test phases, but rerandomized for each subject. The items were presented one at

a time on the screen with the prompt "Is this item grammatical <Y or N>?"
Trials were subject-paced.

Prior to test, subjects were informed that all of the letter strings they had
seen at training had been generated by a complex set of rules which determined
which letters could follow which other letters and which letters could be repeated
and where they could be repeated. They were informed they would see some more
letter strings, only half of which follow the rules, and were to decide whether
they follow the rules or not. It was explained that the complex set of rules were
called an artificial grammar and they had to enter <Y> if they thought an string
obeyed the grammatical rules and <N> if they believed an string violated the
rules. Before the changed letter-set test, the subjects were also informed of the
nature of the letter change, with the use of a hypothetical example to ensure the
process of letter translation was understood. Subjects were then fully debriefed.

At the end of the test phase, subjects were asked how they had approached and
attempted to perform each of the three parts of the experiment. The subjects
were encouraged to elaborate upon any thoughts they had concerning their
performance. Subjects were then fully debriefed.

Results and Discussion

Results of these experiments are summarized in Table 3. Experiment 4, the
incidental repetition condition, was intended to test whether repeated, but
incidental, coding of the surface structural sequences of training items led to a
knowledge base capable of supporting transfer to both same and changed letter-
set classification tests. In the same letter-set test, in which the test items were
instantiated upon the same letter-set as the training items, subjects in the

incidental repetition condition classified novel grammatical and nongrammatical items at a level reliably above chance (mean percent correct 56.3% versus 50% chance level; $t(25) = 4.765$, $p < .001$). However, in the changed letter-set test, in which the same test items were translated into the alternate letter-set, classification performance ($M = 51.5\%$) was not reliably above chance, $t(25) = .906$, $p = .374$, and was reliably inferior to classification performance in the same letter-set test, $t(25) = 2.675$, $p = .013$. This inferior performance in the changed letter-set test was a result of both a lower proportion of correctly classified grammatical items (hits) and a lower proportion of correctly classified nongrammatical items (correct rejections). That is, both the proportion of hits ($p = .588$ versus $p = .559$) and the proportion of correct rejections ($p =. 539$ versus $p = .471$) were higher in the same letter-set test than in the changed letter-set test.

Table 3
Experiments 4-6: Probability of Correct Classification
as a Function of Induction Task and Test Letter Set

| Experiment | Induction Task | Test Letter Set | |
		Same	Changed
4	Incidental Repetition	.56	.52
5	Incidental Analysis	.56	.57
6	Memorization	.59	.53

It was concluded that subjects in the incidental repetition condition gained knowledge about the surface structural regularities of the set of grammatical

training items and could apply this knowledge to the same-set test, in which this form of information is applicable, but could not apply this knowledge to the changed letter-set test, in which this form of information is not applicable. Because the subjects had no motive to pay a great deal of attention to the training stimuli, this surface structural knowledge base was quite likely fragmentary in nature, including only parts of the surface structural sequences that were incidentally repeated in training. Subjects did not appear to acquire knowledge of deep structural regularities which would be useful for the changed letter-set test. This is a result of the subjects having no time or motive to perform any processing upon the deep structural features of the training items, or seek commonalities between patterns expressed by the training items which would result in knowledge pertaining to the deep structural regularities of the set of training items.

The most important conclusion concerning the results of Experiment 4, is that, although the subjects acquired knowledge capable of supporting same letter-set grammatical discrimination, this form of knowledge was not capable of supporting changed letter-set discrimination. This is directly contrary to the structural abstraction account which suggested that the deep, rule-based structural regularities underlying a set of training items are automatically acquired under incidental conditions, and that this form of abstracted knowledge is directly applicable to both same and changed letter-set tests. The results of Experiment 4 suggest that memory does not automatically acquire deep structural knowledge.

Experiment 5, the incidental analysis condition, was designed and conducted to induce subjects to code the deep structural features of the grammatical training items. It was expected that this form of coding would result in a knowledge base capable of supporting both same and changed letter-set transfer. In the same letter-set test subjects in the incidental analysis condition were able to classify novel grammatical and nongrammatical items ($\underline{M}$ = 55.8%) at a level reliably above chance, $\underline{t}(25)$ = 3.377, $\underline{p}$ = .002. Moreover, in the changed letter-set test, classification performance ($\underline{M}$ = 57.1%) was also reliably above chance, $\underline{t}(25)$ = 6.480, $\underline{p}$ < .001. There was no reliable difference between test conditions $\underline{t}(25)$ = .604, $\underline{p}$ = .551. Therefore, subjects had knowledge that was equally applicable to both same and changed letter-set classification tests. This conclusion is reinforced by the fact that in both the same and changed letter-set tests the proportions of hits ($\underline{p}$ = .604 versus $\underline{p}$ = .603) and false alarms ($\underline{p}$ = .512 versus $\underline{p}$ = .535) were highly similar.

It can be concluded that subjects in the incidental analysis condition acquired knowledge of the deep structural regularities underlying the set of grammatical training items. As argued in the general introduction, only deep structural knowledge could support both same and changed letter-set grammatical discrimination. Subjects in the incidental analysis were induced to analyze the deep structural features of the training items. Specifically, the placement of the bar marker mostly likely caused the subjects to break down each training stimulus into separable deep structural features in order to make the necessary repetition detection decision. Therefore, the representations acquired in this condition were most likely separate codes of each of the individual deep

structural features occurring in the grammatical training items, rather than codes of the conjoint occurrence of these deep structural features within a training item found in Experiment 3. However, the exact nature of this representation is difficult to determine. Subjects could have coded these deep structural features in terms of the letters that were actually found to repeat in the training items. Alternatively, they could have coded this information in a more abstract way by simply attending to the location of the repeated letter, but not the identity of the letter that was repeated. Therefore, the surface structural identity of the repeated letters in the training items may or may not have formed part of the representations of training experiences.

It is crucial that the changed-letter grammatical sensitivity demonstrated in Experiment 5 was not likely the result of automatic abstraction of the deep structural regularities across items. This conclusion is reinforced by the results of Experiment 4. In both cases, the particular operations performed by the subjects on item structure had a critical influence on the form of their structural sensitivity. In Experiment 4, the coding of the surface structural sequences of the training items resulted in surface structural knowledge capable of supporting discriminating between grammatical and nongrammatical items based upon the same letter-set. In contrast, in Experiment 5, the coding of the deep structural features of the training items resulted in deep structural knowledge capable of supporting discrimination between grammatical and nongrammatical items based on either letter-set. In both cases, the particular form of processing was applied to satisfy the primary demands of the induction task. Clearly, the knowledge acquired in incidental grammar learning is

conditional, depending upon the particular operations performed on particular aspects of item structure in order to achieve some specific incidental purpose. Therefore, the knowledge acquired does not appear to be the result of automatic structural abstraction.

Experiment 6, the memorization condition, was designed and conducted to determine the form of the knowledge base acquired by subjects who perform under this standard training procedure. Memorization of training stimuli has been shown to support grammatical sensitivity to both same and changed letter-set stimuli (Mathews et al., 1989; Reber, 1969). This was also in the case in this experiment, with classification of novel grammatical items instantiated upon both the same letter-set ($\underline{M}$ = 59.2%) and the changed letter-set ($\underline{M}$ = 52.6%) reliably above chance, respectively, $\underline{t}(25)$ = 5.959, $\underline{p}$ < .001 and $\underline{t}(25)$ = 2.087, p = .047. However, classification performance in the same letter-set test was reliably superior to classification performance in the changed letter-set test, $\underline{t}(25)$ = 3.638. $\underline{p}$ < .001. This indicates that the major effect of memorization was to make available surface structural rather than deep structural knowledge. Moreover, inferior performance in the changed letter-set condition was largely a result of a reduction in the proportion of correctly classified grammatical items (hits) rather than a reduction in the proportion of correctly classified nongrammatical items (correct rejections). That is, the proportion of hits in the same letter-set test ($\underline{p}$ = .652) was much higher than the proportion of hits in the changed letter-set test ($\underline{p}$ = .535), but the proportion of correct rejections was not ($\underline{p}$ = .533 versus $\underline{p}$ = .517, respectively). Therefore, success relied largely upon the correct classification

of grammatical items and this occurred most frequently when test items were instantiated upon the same letter-set as the training items.

In support of the conclusion that memorization made available mostly surface structural knowledge it is necessary to compare the changed letter-set transfer in the memorization condition to the changed letter-set transfer in the incidental analysis condition. The incidental analysis condition should have supported greater changed letter-set transfer because was concluded to support the acquisition of a deep structural knowledge base. In fact, changed letter-set transfer was reliably greater following incidental analysis than following memorization, $t(50) = 2.724$, $p = .009$. Therefore, it was concluded that there was a greater availability of deep structural knowledge following incidental analysis than following memorization, and that memorization, as indicated by good same letter-set transfer, supports the acquisition of mostly surface structural knowledge.

These differences in performance appear to indicate that memorization instructions caused subjects to primarily code the surface structural sequences of the training items. That is, subjects in the memorization condition coded the training items in a way similar to subjects in the incidental repetition condition. The difference between these two conditions was that subjects in the memorization condition, informed that there would be a later memory test, had a direct motive to rehearse this form of information. It is also apparent from the reports given by subjects concerning their own memorization strategies that approximately half of them also attempted, at least once, to analyze the deep structural features of the training items. A tentative conclusion is that this form

of deep structural item coding was responsible for the small amount of changed letter-set transfer demonstrated in the memorization condition. However, because no subject reported consistently analyzing the deep structural patterns of individual training items throughout the entire training phase, it is difficult to determine the effect that this processing strategy had.

Memorization instructions, in fact, leave the experimenter with very little opportunity to determine which of the critical aspects of item coding are responsible for grammatical discrimination. In light of the results of Experiment 4 and 5, in which the well defined encoding task had critical influence on the form of knowledge acquired by subjects, it is clear that memorization instructions do not give the experimenter enough control over the critical aspects of item coding. Subjects given memorization instructions adopt various strategies to encode the training item that range from the form of encoding performed in the incidental repetition condition to the form of encoding performed by subjects in the incidental analysis condition. The idea that under memorization instructions, memory reveals its highly adaptive ability to unconsciously and automatically abstract the complex rule-governed regularities of a set of stimuli is both powerful and exciting (Reber, 1989). However, given the results of the current experiments, it appears far more likely that memorization is simply a task in which subjects perform various encoding operations on the stimuli presented, in highly adaptive preparation for a later memory test whose form has not been specified.

It is concluded that the demands of the induction task caused the processing of different aspects of the structure of individual training items. Sensitivity to the

deep structural regularities of grammatical items was served by the coding of the deep structural features of a set of grammatical training items. Sensitivity to the surface structural regularities of a set of grammatical items was served by the coding of the surface structural sequences of grammatical training items. In both cases, it is thought that the operations performed to satisfy the primary demands of the induction task resulted in representations of different aspects of the structure of individual training items, which served to later support a specific form of sensitivity to set-wise structural regularities. It is concluded that there is no reason, given the evidence, to posit that memory has additional processes specially adapted to automatically abstract structure across grammatical training items.

General Discussion

As predicted by the episodic-processing account, the purpose for which subjects incidentally encountered grammatical training items was shown to have a critical impact on later classification of novel grammatical and nongrammatical items. In Experiment 4, it was demonstrated that subjects who were led to code the surface structural sequences of the training items could later discriminate between grammatical and nongrammatical items instantiated upon the same letter-set, but not between the same items based on a different letter set. In contrast, Experiment 5 demonstrated that subjects who were led to code the deep structural aspects of training items could later discriminate between grammatical and nongrammatical items presented in either the same or a changed letter-set. This indicates that the operations performed to satisfy the primary

demands of the induction task caused memory to form representations of different aspects of the structure of individual training items. These representations later supported different but specific forms of sensitivity to set-wise structural regularities.

Experiment 3 was conducted to investigate whether changed letter-set transfer is supported by particular knowledge of individual items, or general knowledge of set-wise structural regularities. It demonstrated that coding the set of deep structural features of individual training items can support identification of items instantiated upon a different letter-set. This indicates that the abstraction of set-wise structural regularities is not necessary for strong transfer to items instantiated on a changed letter-set. Together, the results of Experiments 3 and 5 demonstrated that subjects can code either the individual deep structural features of individual items, or compounds of those deep structural features depending upon the nature of the task demands.

These results support the episodic processing account's contentions that: 1) in order to satisfy task demands, memory processes particular aspects of training item structure and forms representations of these processing experiences; and 2) grammatical sensitivity is dependent upon the extent to which these representations of past processing experiences are applicable to specific tests of grammatical knowledge.

With respect to the first point, the variable and specific nature of the knowledge acquired in implicit learning appears to result from memory's ability to process specific aspects of stimulus structure for particular purposes. Every purpose given to a subject invokes operations that are functional in achieving

that purpose. For example, subjects can process the surface structural sequences of individual training items in order to achieve their intended purpose of repeating the letters of the string out loud. Alternatively, subjects can process the deep structural features of individual training items in order to achieve their intended purpose of judging whether individual letters are repeated or not repeated. Memory appears to form representations of these particular processing experiences. Moreover, with respect to the second point, grammatical discrimination appears to be supported only to the extent that the processing performed at test can make use of the products of the specific processes performed in training. For example, the products of surface structural coding can support transfer to items based upon the same letter-set, but not items based upon the changed letter-set. This is possible because items instantiated upon the same letter-set share surface structural sequences, but items instantiated upon different letter-sets do not. In contrast, representations of the deep structural features of individual training items can support transfer to both the same and changed letter-set classification tests. This is possible because grammatical items, regardless of the letter-set they are instantiated on, share deep structural features.

It is concluded that the current results, and implicit learning in general, can only be explained by a theory that claims that memory processes different aspects of stimulus structure under the guidance of an incidental purpose given by the experimenter, forms a representation of these processing experiences, and retrieves these experiences to support later performance under similar

circumstances. Therefore, the episodic-processing account is the only current account that can fully explain implicit learning.

Reber's Structural Abstraction Account and the Current Experiments

The results of the current experiments cannot be explained by Reber's (1989) theory of implicit grammar learning, which posited a default tendency to abstract structural regularities. The results of Experiment 5 demonstrated that sensitivity to deep structural regularities is conditional upon the coding of the deep structural features of individual training items. Moreover, Experiment 4 demonstrated that subjects can become sensitive to set-wise surface structural regularities, without becoming sensitive to general deep structural regularities. These results cannot be easily explained by the automatic abstraction of deep structural regularities across experience or a switch to an instance mode of learning. These results indicate that sensitivity to set-wise deep and surface structural regularities is conditional on performing operations upon particular aspects of item structure. Such conditional sensitivity is not part of Reber's account of implicit learning. In addition, Experiment 3 demonstrated that subjects can code the conjoint occurrence of deep structural features of individual items in order to support transfer across letter-sets. This demonstrated that changed letter-set transfer can be supported by knowledge about individual items. Given that changed letter-set transfer was considered by Reber (1989) to be the strongest evidence for set-wise structural abstraction, this result places Reber's theory of implicit learning in further doubt. Therefore, Reber's structural abstraction account cannot explain the current results, whereas, the episodic-processing account can.

There is one aspect of the Reber's account of implicit learning that is superficially similar to the episodic-processing account. This is the contention that the nature of the processing performed by subjects given different task requirements is functional. Reber (1989) suggested that the learner "assumes a cognitive stance that is functional, that will accomplish the task at hand" (p. 226). Reber argued that under observation or standard memorization training conditions it is functional to abstract the structural regularities underlying a set of items. When the subject is strongly encouraged to code instances it is more functional to switch to an instance learning mode. However, it is unclear why it is only functional to code instances when performing paired-associate learning and not standard memorization. In both cases the primary requirement of the task is to memorize items for a later test. Therefore, it should be equally functional to switch to an instance learning mode in both cases. Moreover, it is difficult to understand why, when observing items, it is functional to abstract general structure. To perform the task of observing items it is only necessary to pay attention to the items, not abstract general structure. Therefore, Reber is presenting a theory in which memory has a default tendency to abstract structure. This default tendency is disrupted when instance coding takes place. As a result, Reber does not appear to be providing a theory in which different modes of learning are functionally applied to the current situation.

Is Artificial Grammar Learning Implicit?

The acquisition and utilization of knowledge in the current experiments appears to be quite complex. The situation facing the subjects in these experiments is somewhat like the following situation. A person has been in a

friend's living room many times. She has walked through the room avoiding the furniture, has sat on the furniture and has looked about while talking to her friend. These everyday purposes have created experiences concerning the arrangement of the furniture in the room. However, the position of the furniture has always been incidental to the primary purpose for encountering the furniture. Moreover, this knowledge has been acquired without anticipation of ever needing knowledge about the position of the furniture. Nonetheless, one day she walks into the room and has the feeling of uncertainty, a feeling that something is different. This feeling could serve as the basis of a conscious decision, namely "that something is different". Moreover, it could compel her to start looking at the room to consciously seek what has changed. Upon looking around the room, she gets a strong feeling that a chair seems to be in a different place. Although she cannot consciously recall where the chair was before, she could state with a great deal of certainty that it had been moved.

Although this example pertains to a specific room, it can also apply to the category of any one type of room. For example, walking into a living room for the first time one can get the feeling that the furniture is arranged in a pleasing manner or a non-pleasing manner. This type of feeling state likely stems from multiple past experiences with rooms of that type, which share commonalities in furniture arrangement. People are implicitly sensitive to the regularities in furniture arrangement. As in the above example, this form of sensitivity could be supported by specific representations of particular experiences. That is, past encounters with the objects in similar rooms for specific purpose could result in representations of those experiences. In turn, these representations could, as in

the case of artificial grammar learning, support sensitivity to general structural regularities.

In the current experiments, performance in the incidental training conditions (Experiments 1, 2, 4 and 5) could be considered analogous to the above example. Subjects in the incidental training conditions, while they were aware they were encountering the training items for a particular purpose, did not anticipate that their encounters with the training stimuli would be relevant to performance on a later test. Because the later test was unanticipated, subjects could not simply attempt to acquire the form of knowledge that they believed would be most applicable to the test task. Therefore, the acquisition of knowledge relevant to performance on a later classification test was incidental. The incidental nature of the learning process made it difficult for the subjects to consciously evaluate what they had learned.

Also analogous to the above example, subjects in the current experiments appeared to have the feeling that something was "different" or "wrong" about nongrammatical test items and that something was the "same" or "right" about grammatical test items. With respect to consciousness, subjects reported having few or no rules or other forms of conscious knowledge with which to classify test items. However, subjects reported consciously attempting to orient their attention toward certain structural characteristics of the test strings in an attempt to determine what is "wrong" or "right" about that item. For example, subjects reported directing their attention toward the pattern of repetition within a test item in the hopes that the whole item, or parts of the item, would feel familiar or unfamiliar. Thus, above-chance classification performance

appears to be largely a result of unconscious processing resulting in a consciously available feeling state. Consciousness appears able to focus attention on different aspects of item structure in order to make that feeling state available.

Thus, the acquisition and utilization of the knowledge supporting grammatical sensitivity does not appear to be entirely independent of the subjects' intentions and strategies. However, because the knowledge incidentally acquired in the current experiments was not easily retrieved into consciousness, the learning process can be described as implicit. This does not mean that subjects were unaware that they were performing operations upon the training stimuli, or that they were unaware that they had knowledge relevant to performing the test task. It only means that the products of their learning were difficult to retrieve to consciousness because such direct retrieval was not supported by the cuing conditions offered in the test situation.

Even if subjects could have retrieved the products of their learning into consciousness, it would not necessarily indicate that classification decisions were made explicitly. The conscious utilization of directly retrieved past experiences may not be capable of supporting grammatical discrimination. That is, retrieving one or two past exemplars into consciousness and comparing them to the current test stimuli may be an inefficient basis for classification response. It is possible that only the reliance upon consciously available feeling states, resulting from implicit parallel access to multiple past experiences, can efficiently support grammatical discrimination. Conscious and unconscious memory processes may use the same products of experience, but in qualitatively

different ways, for qualitatively different purposes (Jacoby & Witherspoon, 1982).

The Interaction Between Processing and Structure

The episodic-processing account suggests that the operations performed on a stimulus are both supported and limited by the affordances of that stimulus' structure. In the current experiments, the structure of the items supported the forms of processing required, but under other circumstances these forms of processing would not have been possible. For example, if the items had been created with lights of different frequencies as elements, then the verbal repetition of surface structural features in the incidental repetition condition would not have been possible. Regehr and Brooks (in press) provided important results and arguments concerning the interaction between stimulus structure and various perceptual operations that can be performed on that structure. Regehr and Brooks showed that the same general analytic structure (set of rules) could be represented by stimuli in various perceptual forms. They demonstrated that the perceptual form of the stimuli, independent of the analytic structure, controlled and limited the form of processing that could be applied to the stimuli. In turn, the form of processing applied to the structure of individual stimuli controlled sensitivity to the underlying analytic structure. Therefore, they concluded that the perceptual manifestation of analytic structure has a critical impact on sensitivity to set-wise analytic structure.

This same impact could be demonstrated in the artificial grammar learning paradigm by manipulating the type of elements used to instantiate the stimuli. For example, a grammatical pattern could be represented by a series of light

flashes in various locations on a computer monitor. That is, each element of a stimulus would be represented by a different location, rather than a different letter. Subjects might be induced to code the individual elements of a grammatical training item in a separable fashion if the series of light flashes were far apart in either distance or time. However, they might code the individual patterns in a more integral fashion if the light flashes were close together in distance or time. This change in the perceptual manifestation of the training stimuli could have a critical effect on later sensitivity to the general structure of the grammar. As indicated by Experiment 1 and 2, the separable coding of individual elements could result in only a small amount of grammatical discrimination. However, in this case, separable element coding would result from the perceptual characteristics of the stimuli, rather than from an overt processing demand.

Future investigation should determine the impact of the perceptual manifestations of the training stimuli in the artificial grammar learning paradigm. However, this investigation should not ignore the fact that the structural properties of stimuli will be experienced, and represented, in terms of the operations a subject performs on the stimuli to achieve a specific purpose.

Implications for Future Research

The current experiments demonstrated the importance of tightly controlling the coding operations performed by subjects when learning. For example, in Experiment 4 and 5, control over the induction task resulted in an ability to determine the critical aspects of item coding that were responsible for grammatical discrimination on the same and changed letter-set tests. This is in

contrast to the more usual memorization instructions, which leave the experimenter with very little opportunity to determine the aspects of item coding that are responsible for grammatical discrimination. Moreover, a tight control over the structure of the set of items was also necessary. For example, in Experiment 3, the nonpredictive nature of any set-wise structural regularities in the set of stimuli made it possible to conclude that the abstraction of set-wise structural regularities is not necessary for strong transfer to items instantiated on a changed letter-set. Even more highly specific control over both the structure of the set of items and the training and testing conditions will be necessary in future investigation of implicit learning.

Future research should include the investigation of a number of partially independent dimensions which could be used to describe the nature of the implicit learning process. For example, the utilization and acquisition of grammatical knowledge could be described as deliberate or incidental, conscious or unconscious, and could be described as occurring with or without awareness. Moreover, the knowledge base acquired could be verbalizable or nonverbalizable. To date, some investigators of artificial grammar learning have assumed that because the knowledge acquired in artificial grammar learning is difficult to verbalize, it must be qualitatively different from consciously available, verbalizable knowledge (e.g., Reber, 1989; see also Berry and Broadbent, 1984, for similar arguments concerning knowledge in other types of complex rule learning situations). In addition, researchers have assumed that the process involved in acquiring verbalizable knowledge is qualitatively different from the process involved in acquiring nonverbalizable knowledge (Dienes et al., 1991;

Reber, 1976, 1989). As a result, these researchers have assumed that people can deploy a deliberate, conscious mode of learning which results in verbalizable knowledge, or an incidental, unconscious mode of learning which results in nonverbalizable knowledge (Berry and Broadbent, 1984; Dienes et al., 1991; Reber, 1976, 1989).

However, the knowledge acquired under incidental and deliberate learning conditions may not be qualitatively different. For example, encouraging subjects to deliberately search for the rules underlying the set of grammatical training exemplars was demonstrated to have no influence upon their later classification performance, or ability to verbalize the knowledge they had acquired (Dienes et al., 1991; Mathews et al., 1989; Perruchet & Pacteau, 1990). This suggests that it is unlikely that there are two neatly defined, stable modes of learning which can be specified by using the dimensions conscious/unconscious, incidental/deliberate and verbalizable/nonverbalizable. It is far more likely that in any situation the acquisition and utilization of knowledge can be described as occurring at some point on each of these partially independent dimensions. The possibility that changes in the learning process are caused by changes along any of these dimensions can only be discovered by independently manipulating particular aspects of the learning situation. For example, subjects could be informed that the training stimuli are created by a set of rules which they must discover, and then given the item analysis training used in Experiment 5. Any differences in classification performance between this deliberate condition and the incidental condition used in Experiment 5 would be the result of a switch from an incidental to a deliberate form of learning. Moreover, the manipulation

of factors that might selectively affect conscious or unconscious memory processes could be used to determine the importance of conscious and unconscious aspects of performance. Manipulations of this nature could result in a better description and understanding of the artificial grammar learning process.

Conclusion

As predicted by the episodic-processing account, the results of the current experiments demonstrated that different purposes for encountering exemplars of a grammar cause critically different operations to be performed on particular aspects of stimulus structure. In the experiments conducted, manipulation of the demands of the task in which exemplars were incidentally encountered and encoded was shown to radically modify the nature of the knowledge base acquired in implicit learning. These modifications in the nature of the knowledge base acquired in implicit learning cannot be explained by the automatic abstraction of structure. However, they can be explained by the episodic-processing account which postulates that memory performs whatever operations are functional in satisfying current demands, preserves the products of those operations, and applies those products in a later test. To date, the episodic-processing account has also had success in explaining phenomena as widely divergent as the word superiority effect (Whittlesea & Brooks, 1988), pseudoword perception (Whittlesea & Cantwell, 1987), semantic priming (Whittlesea & Jacoby, 1990), illusions of memory driven by perceptual and conceptual fluency (Whittlesea, 1992; Whittlesea, Jacoby & Girard, 1990) and sensitivity to typicality in clustered categories (Whittlesea, 1987). Moreover, similar processing accounts have been used to explain performance on explicit and

implicit measures of memory (e.g., Jacoby, 1983; Masson & MacLeod, 1992; Roediger, 1990; Roediger, Weldon & Challis, 1989). The episodic-processing account promises to provide a general framework of mental performance capable of explaining remembering, attention, perception and learning.

References

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

Anderson, J.R. (1978). Arguments concerning representations for mental imagery. Psychological Review, 85, 249-279.

Anderson, J.R. (1979). Further arguments concerning representations for mental imagery: A response to Hayes-Roth and Pylyshyn. Psychological Review, 86, 395-406.

Berry, D. C. & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. Quarterly Journal of Experimental Psychology, 35, 209-231.

Brody, N. (1989). Unconscious learning of rules: Comment on Reber's analysis of implicit learning. Journal of Experimental Psychology: General, 118, 236-238.

Brooks, L. R. (1978). Non-analytic concept formation and memory for instances. In E. Rosch and B. B. Lloyd (Eds.), Cognition and categorization (pp. 169-211). Hillsdale, NJ: Erlbaum.

Brooks, L. R., & Higham, P. A. (1992). Learning the experimenter's design: Tacit sensitivity to the structure of memory lists. Submitted.

Brooks, L. R., & Vokey, J. R. (1991). Abstract analogies and abstracted grammars: Comments on Reber (1989) and Mathews et al. (1989). Journal of Experimental Psychology: General, 120, 316-323.

Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. Journal of Experimental Psychology: General, 104, 268-294.

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

Druhan, B.B., & Mathews, R. C. (1989). THIYOS: A classifier system model of implicit knowledge of artificial grammars. Paper presented at the 11th annual meeting of the Cognitive Science Society, Ann Arbor.

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.

Estes, W. K. (1986). Memory storage and retrieval processes in category learning. Journal of Experimental Psychology: General, 115, 155-174.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Psychology, 18, 155-170.

Glenberg, A., Smith, S. M., & Green, C. (1977). Type 1 rehearsal: Maintenance and more. Journal of Verbal Learning and Verbal Behavior, 16, 339-362.

Gordon, P. C., & Holyoak, K. J. (1983). Implicit learning and generalization of the "mere exposure" effect. Journal of Personality and Social Psychology, 45, 492-500.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. Psychological Review, 93, 411-428.

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

Jacoby, L. L. (1983). Remembering the data: Analyzing interactive processes in reading. Journal of Verbal Learning and Verbal Behavior, 22, 485-508.

Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception, and concept formation. In G. H. Bower (Ed.), The Psychology of Learning and Motivation (Vol. 18, pp. 1-47). New York: Academic Press.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. Journal of Experimental Psychology: General, 103, 306-340.

Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. Journal of Experimental Psychology: General, 118, 126-135.

Jacoby, L. L., & Witherspoon, D. E (1982). Remembering without awareness. Canadian Journal of Psychology, 32, 300-324.

Kolers, P. A. (1973). Remembering operations. Memory & Cognition, 1, 347-355.

Kolers, P. A. (1979). A pattern-analyzing basis of recognition. In L. S. Cermak & F. I. M. Craik (Eds.), Levels of processing in human memory. (pp. 363-384). Hillsdale, NJ: Erlbaum.

Kolers, P. A., & Roediger, H. L. (1984). Procedures of mind. Journal of Verbal Learning and Verbal Behavior, 23, 425-449.

Koiers, P. A., & Smythe, W.E. (1984). Symbol manipulation: Alternatives to the
computational view of mind. <u>Journal of Verbal Learning and Verbal Behavior,</u>
<u>23,</u> 289-314.

Kushner, M., Cleermans, A., & Reber, A. (1991). Implicit detection of event
interdependencies and a PDP model of the process. Paper presented at the
12th annual conference of the Cognitive Science Society, Chicago.

Mandler, G. (1980). Recognizing: The judgment of previous occurrence.
<u>Psychological Review, 87,</u> 252-271

Martin, E. (1971). Verbal learning and the independent retrieval phenomenon.
<u>Psychological Review, 78,</u> 314-332.

Masson, M. E. J., & MacLeod, C. M. (1992). Reenacting the route to
interpretation: Enhanced perceptual identification without prior perception.
<u>Journal of Experimental Psychology: General, 121,</u> 145-176.

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

Mathews, R. C. (1991). The forgetting algorithm: How fragmentary knowledge of
exemplars can abstract knowledge. <u>Journal of Experimental Psychology:</u>
<u>General, 120,</u> 117-119.

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. <u>Journal of Experimental Psychology:</u>
<u>Learning, Memory and Cognition, 15,</u> 1083-1100.

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

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. Psychological Review, 85, 207-238.

Morris, C. C., Bransford, J. D., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. Journal of Verbal Learning and Verbal Behavior, 16, 519-533.

Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition and typicality. Journal of Experimental Psychology: Learning, Memory and Cognition, 14, 700-708.

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

Reber, A. S. (1967). Implicit learning of artificial grammars. Journal of Verbal Learning 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. (1989). Implicit learning and tacit knowledge. Journal of Experimental Psychology: General, 118, 219-235.

Reber, A. S., & Allen, R. (1978). Analogic and abstraction strategies in synthetic grammar learning: A functionalist interpretation. Cognition, 6, 193-221.

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 a the form and structure of a body of tacit knowledge. Cognition, 5, 333-361.

Regehr, G., & Brooks, L.R. (in press). Perceptual manifestations of an analytic structure: The priority of holistic individuation. Journal of Experimental Psychology: General.

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

Roediger, H. L., Weldon, M. S., & Challis, B. H. (1989). Explaining dissociations between implicit and explicit measures of retention: A processing account. In H. L. Roediger and F. I. M. Craik (Eds.), Varieties of memory and consciousness: Essays in honor of Endel Tulving (pp. 3-41). Hillsdale, NJ: Erlbaum.

Roussel, L.G., Mathews, R.C., & Druhan, B.B. (1990). Rule induction and interference in the absence of feedback: A classifier system model. Paper presented at the 12th annual meeting of the Cognitive Science Society, Chicago.

Rundus, D. (1977). Maintenance rehearsal and single-level processing. Journal of Verbal Learning and Verbal Behavior, 16, 665-681.

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

Tulving, E., & Thompson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. Psychological Review, 80, 352-373.

Vokey, J. R., & Brooks, L. R. (1992). The salience of item knowledge in learning artificial grammars. Journal of Experimental Psychology: Learning, Memory & Cognition, 18, 328-344.

Whittlesea, B. W. A. (1983). The representation of concepts: An evaluation of the abstractive and episodic perspectives. Unpublished doctoral thesis, McMaster University, Hamilton, Ontario.

Whittlesea, B. W. A. (1987). Preservation of specific experiences in the representation of general knowledge. Journal of Experimental Psychology: Learning, Memory and Cognition, 13, 3-17.

Whittlesea, B. W. A. (1992). Illusions of familiarity. Manuscript in preparation.

Whittlesea, B. W. A., & Brooks, L. R. (1988). Critical influence of particular experiences in the perception of letters, words and phrases. Memory & Cognition, 16, 387-399.

Whittlesea, B. W. A., Brooks, L. R., & Westcott, C. (1991). Utilization of analytic and nonanalytic processing: Factors affecting reliance on general and particular knowledge. Manuscript in preparation.

Whittlesea, B. W. A., & Cantwell. A. L. (1987). Enduring influence of the purpose of experiences: Encoding-retrieval interactions in word and pseudoword perception. Memory & Cognition, 15, 465-472.

Whittlesea, B. W. A., & Dorken, M. D. (1992). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. Manuscript submitted for publication.

Whittlesea, B. W. A., & Jacoby, L. L. (1990). Interaction of prime repetition with visual degradation: Is priming a retrieval phenomenon? Journal of Memory and Language, 29, 546-565.

Whittlesea, B. W. A., Jacoby, L. L., & Girard, K. (1990). Illusion of immediate memory: Evidence for an attributional basis of familiarity and perceptual quality. Journal of Memory and Language, 29, 716-732.

Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. American Psychologist, 35, 151-175

Appendix 1

Training and Test Items of Experiment 1 and 2

Training Grammatical	Test Grammatical	Test Nongrammatical
MTTTC	CXTTTC	CXRRT
MTTCT	MTTTC	CXX
MTC	MTTCRX	CXRCM
MTCRX	MCRXCT	XCRXRR
MTCRXM	MTCRXC	XTTTTC
MCRX	MTCRXR	MTCC
MCRXRR	MCRXM	MMCRX
MCRXT	CXCRXR	MCRTR
MCRXC	MTTTCT	MTRCRX
MCRXCT	CXRM	TTCT
CXM	MCT	MTTCTR
CXRR	MTCT	TCTTXC
CXRRM	MTTC	RCT
CXRRRR	MCRXR	MXCT
CXTTCT	CXRRR	CRRRM
CXTCRX	CXTC	XRCXC
CXTCT	CXR	CCXRM
CXCRX	CXCT	CXRT
CXCRXC	MTC	MTRC
CXCT	CXRRRM	CXMRXC
	CXTTC	MTM
	CXC	TXRRM
	CXCRX	MXCRXM
	CXCRXC	MTCRTR
	MCRXRM	RRRXC

Appendix 2

Training and Test Items of Experiment 3

Training Patterns	Training Items	Old Test Items	New Test Patterns	New Test Items
1234-1324	HDFX-HFDX	GZTP-GTZP	1234-1243	LKPG-LKGP
1234-1342	DMXH-DXHM	GTLP-GLPT	1234-4312	VCGT-TGVC
1234-2413	SMFB-MBSF	PTZC-TCPZ	1234-3421	GKTL-TLKG
1234-2431	SDRF-DFRS	CVZL-VLZC	1234-2134	VKTG-KVTG
1234-3124	DHMR-MDHR	VLCK-CVLK	1234-1423	TCZP-TPCZ
1234-3142	XHMF-MXFH	CKVT-VCTK	1234-2314	VTCK-TCVK
1234-4231	XSMB-BSMX	PLZT-TLZP	1234-3241	ZPTG-TPGZ
1234-4213	SRDF-FRSD	TKLV-VKTL	1234-4132	PLGT-TPGL
1123-3244	BBHX-XHRR	ZZLC-CLPP	1123-4432	LLCP-KKPC
1123-3442	FFRD-DSSR	PPZT-TLLZ	1233-2144	TCZZ-CTGG
1223-4134	MHHR-SMRS	VKKZ-GVZG	1231-4324	LGKL-ZKGZ
1223-4431	DMMR-SSRD	CPPK-VVKC	1123-4234	CCTL-VTLV
1231-2344	SDFS-DFHH	LVPL-VPGG	1223-3144	ZCCV-VZTT
1231-2443	BSHB-SDDH	PGTP-GVVT	1233-2441	GTLL-TPPG
1233-4421	RBFF-XXBR	GTKK-VVTG	1231-4423	CPKC-LLPK
1233-4124	DXBB-HDXH	ZPGG-CZPC	1223-3441	GTTL-LPPG

Note: Training and test items listed are examples only. Actual stimuli varied between subjects: See text.

Appendix 3

Training and Test Items of Experiment 4-6

Training Items	Same-Set Test Items		Example of Different-Set Test Items	
Grammatical	Grammatical	NonGrammatical	Grammatical	NonGrammatical
RXXCTXT	RXXCRCT	RCTCTXT	FPPLFLS	FLKLKSK
RTRCCRX	RXTTXXT	RXRCMCT	FPLLPPL	PSPKLKF
RXTTMCT	RTTXTRX	RTCXMCT	LFFKFLK	LPKSFKP
RTTXTCT	RTTXMCT	RTCCXXT	SLLFKPL	KPLLSSP
RTTXCRX	RTCXXMX	RMCCXCT	PLSFFKF	FLSSPSK
RTCXTRX	RTRMTRX	RTRTRXT	SPSKPSF	FLFLFSL
RTRMXXT	RTRMTCT	RXTXXMX	LSLFSKS	LKSKFPK
RMCCTMX	RMCCTXT	RTCMTXX	KSFFLPL	SPKLPFF
RMCCRCT	RMCXTCT	RTCCTMX	KPSLFSF	LFPPFSK
RMCXTRX	RMCXXMX	RMRTRTX	FKPSSKS	KLKPKPF
RMRTXXT	RMRTMCT	RRCXTRX	PFPSFLS	SSPLKSL
RMRTXCX	MXMRTMX	MTRXTRX	SKSFLSK	SPKLPKL
MXMRTXT	MXTTXCX	MRTTXTX	PFSSFKF	KLSSFSF
MXTTXMX	MTTXXCX	MTCXRRX	KFFLLPL	PLSFKKF
MTTXTRX	MTRMXXT	MTCTTMX	PSLPKKS	SLPLLSF
MTRMTRX	MTCXTRX	MCCMTCT	FSKLSPL	PSSPLSL
MTRMTCT	MXCRTXT	MXMRTRT	FPSKLPL	PKPFSFS
MTCXXMX	MXCRTMX	MCCMXXT	SPFKLSP	LPPLFFS

Appendix 3 (continued)

Training Items	Same-Set Test Items		Example of Different-Set Test Items	
Grammatical	Grammatical	NonGrammatical	Grammatical	NonGrammatical
MTCXTCT	MXCMTRX	MXCXCRX	KLPKSFL	SPFPFKP
MXCRCCX	MXCTMCT	MTRCTRX	KFSPKSP	PKLSKLF
MXCMXXT				
MXCMTCT				
MXCTXMX				
MXCTXXT				
MTTMRCT				

Note: The changed-letter test items presented are only examples. In practice, letters were randomly assigned to the deep structure of the items on a trial-by-trial basis.