

Rule-based and exemplar-based classification in artificial grammar learning

MARY PAT McANDREWS and MORRIS MOSCOVITCH
Erindale College, University of Toronto, Mississauga, Ontario, Canada

In this study, we examined the induction of syntactic rules, given the presentation of letter strings generated from a finite-state grammar. Our primary interest was whether application of abstracted syntax or analogy to remembered exemplars could serve as a basis for judgments of grammaticality of novel stimuli. The grammatical status of test items and their objective similarity to studied exemplars were manipulated independently to investigate whether rule-based or instance-based information was a more important determinant of classification performance. When group data were examined, the results indicated that both factors were equally important in influencing grammaticality judgments about novel letter strings. There were, however, large individual differences in the magnitude of grammatical status effects, with a subgroup of subjects clearly using a classification strategy other than analogy to remembered exemplars. The results offer qualified support for the hypothesis (Reber & Allen, 1978) that rule-based information can be implicitly abstracted given limited experience with richly structured stimulus domains, and these results are inconsistent with a strong version of the instance-based model of categorization.

Recent attempts to specify concept formation and categorization processes have focused on the learning of complex rule structures and ill-defined categories which have no immediately apparent criterial features. Judgments of well-formedness or category membership of items in these complex domains cannot be based on any single easily specified rule or list of defining attributes. In general, acquisition of the objective category or rule relations is investigated by allowing subjects limited experience with exemplars and assessing classification accuracy on a transfer test using novel stimuli. These experiments form the empirical basis for two broad classes of concept formation and categorization models which can be distinguished by the information-processing operations assumed to underlie category-membership judgments.

According to one type of model, subjects categorize novel instances on the basis of some memorial representation of the patterns of invariance inherent in the exemplar set. This model presumes that individuals abstract structural or featural information from the training exemplars and store this information as a high-level generalization, such as a conceptual schema or prototype, which qualitatively differs from a compendium of individual exemplar memory traces. Several researchers have provided support for the prototype model (e.g., Homa & Vosburgh, 1976; Strange, Kenney, Kessel, & Jenkins, 1970) and for other models of classification by abstraction (e.g., Hayes-Roth & Hayes-Roth, 1977; Reed, 1972). For example, Posner and Keele (1968, 1970) reported that category pro-

types that had never been seen were classified with greater accuracy than other new items and that, after a delay, the probability of recognizing these prototypes was greater than for the actual exemplars on which subjects had been trained.

The alternative models regard classification as exemplar-based. That is, decisions about category membership of novel items are based solely on retrieval of information about specific exemplars rather than more abstract information (e.g., Medin & Schaffer, 1978; Medin & Smith, 1981). Thus, whether or not the subject has been able to abstract any higher level information, his classification judgments are guided by similarity of the test item to stored exemplars. In support of this model, Hintzman and Ludlum (1980) replicated the Posner and Keele (1970) findings of differential forgetting of prototype and instance information using a computer simulation of classification based solely on stored exemplars.

Several studies have addressed this issue by examining acquisition and use of syntactical information embodied in letter strings generated from a finite-state artificial grammar (e.g., Reber & Allen, 1978; Vokey & Brooks, in press). Such a grammar, as shown in Figure 1, consists of a vocabulary set (letters M, V, T, R, X) and a set of transition rules which designate the appropriate letter sequences for sentence formation. A sentence (e.g., MTVX) is made up of the letters produced by a set of permissible transitions from the initial state (1) to an exit state (5, 6, or 7). The artificial language is simply all sentences that can be so generated. As is the case with grammars of natural language systems, rule violations yield nongrammatical sentences (e.g., VXV). Although such grammars, like their natural-language counterparts, can be characterized by a finite set of rules defining category membership (i.e., grammaticality), the learning situation

This research was supported by grants to Morris Moscovitch from the Ontario Mental Health Foundation and the Medical Research Council of Canada and a special research program grant from the Connaught Fund, University of Toronto. Requests for reprints should be sent to M. P. McAndrews, Department of Psychology, Erindale College, University of Toronto, Mississauga, Ontario, Canada L5L 1C6.




Figure 1. Finite-state diagram of grammar used in the experiment.

is analogous to other research on ill-defined categories: The subject is confronted with a complex category with a potential infinity of exemplars for which no readily apparent set of relations governs membership and exclusion.

Reber and his colleagues have used this paradigm to investigate abstraction in an extended series of studies (Allen & Reber, 1980; Reber, 1969, 1976; Reber & Allen, 1978; Reber, Kassin, Lewis, & Cantor, 1980). They report that subjects who have studied grammatical letter strings, without knowing about the rule-governed nature of the stimuli, can subsequently classify new grammatical and nongrammatical items with greater than chance accuracy, although they may be unable to articulate specific rules guiding their decisions. Reber interprets these results as evidence for an abstraction model of classification. He argues that subjects form an abstract memorial representation of the structural principles through nonconscious induction of the patterns of invariance embodied in the exemplar set. This veridical, but incomplete, representation is the basis for judging the grammaticality of novel letter strings.

Reber argues further that such performance reflects an implicit acquisition process rather than the use of conscious analytic strategies (cf. Reber, Allen, & Regan, 1985). Empirical support for this claim is provided in several experiments on the effect of variation in training instructions on subsequent classification performance. In one study (Reber, 1976), classification accuracy for subjects instructed to look for underlying rules in exemplars during the presentation phase was below that of naive subjects, and in many instances below chance. Such poor performance was presumed to reflect the generation and use of incorrect hypotheses which interfered with the implicit abstraction process. However, instructions requiring explicit rule induction during study can aid subsequent classification accuracy when exemplars are presented in a manner that makes the patterns of invariance highly salient, such as blocking items according to grammatical subrules (Reber et al., 1980). In addition, classification by abstraction is thought to be disrupted when exemplar presentation takes place in the context of a task, such as paired-associate learning, which focuses all attentional resources on exemplar memorization and induces a conscious exemplar-based classification strategy (Reber & Allen, 1978).

Brooks (1978; Vokey & Brooks, in press) attempted to demonstrate that Reber's findings could be accommodated by an exemplar-based model of concept formation.

Arguing that grammaticality and similarity to exemplars were confounded in Reber's studies, Vokey and Brooks (in press) manipulated the factors orthogonally in the construction of test items. Thus, all transfer stimuli were defined by a combination of similarity distance from an individual studied exemplar and grammatical status. Vokey and Brooks reported that similarity distance was a much stronger determinant of subjects' classification judgments than was grammatical status. In one study, for example, the magnitude of variance in "grammatical" responses accounted for by similarity was 12 times greater than that attributable to grammaticality. Brooks argues that during the initial presentation of exemplars, subjects are forming a memory set consisting of representations of individual items, rather than acquiring some abstract representation of the grammar. Subsequent classification of novel test items is based on an analogy strategy, in which test items serve as retrieval probes to activate the representation of the most similar item in the memory set. Thus, if a test item is highly similar to one of the previously viewed exemplars, it is accepted as grammatical; if no similar exemplars are retrieved, the new item is rejected.

Vokey and Brooks (in press) discuss the effects of manipulating learning tasks and instructions in light of parameters assumed to affect the success of an analogy strategy. Variations in exemplar encoding are said to affect the degree of individuation of the memory-set items, which in turn influences how effectively subjects can exploit similarities between training and transfer items. Thus, under encoding conditions in which subjects are presumed to learn little about each of the training items, such as simple observation, many transfer items may seem vaguely similar to the stored exemplars, and therefore may be labeled grammatical. Brooks argues that these changes in individuation and breadth of transfer, rather than changes in the efficacy of an implicit abstraction process, account for differences in classification performance under different instructional sets and training procedures.

Both Reber and Brooks readily acknowledge the importance of training tasks and instructions in determining the nature of the information a subject has available for classification decisions. The acquisition strategy adopted may be one of explicit rule induction, implicit abstraction, or memorization of individual items, and the classification decisions may be based on any or all sources of information. Other researchers (e.g., Elio & Anderson, 1981; Homa, Sterling, & Trepel, 1981) have also stressed the need for a more functionalist interpretation of categorization performance. The major points of difference between Brooks and Reber are whether more weight should be given to abstraction or analogy as the principal mode of classification, and the extent to which implicit abstraction, as characterized by Reber, is a viable explanatory construct.

The present experiment, which is an extension of Vokey and Brooks's study (in press, Experiment 1), was designed to eliminate some ambiguity regarding the viabil-

ity of an abstraction interpretation of performance in the classification task. We felt that certain aspects of their experimental design introduced a strong bias against potential evidence for implicit abstraction. Contrary to Brooks's arguments, our hypothesis was that, even with the effects of similarity distance and grammatical status separated, an abstraction process could be seen to account for classification performance when several aspects of the learning situation were changed.

First, we presented twice as many exemplars at study with fewer repetitions of each item, because we felt that Vokey and Brooks's presentation design may have induced overlearning of individual exemplars and thereby optimized an analogy strategy in later classification. Even if it were possible to abstract rules in this situation, most individuals, confident they know nothing about rules to which they had not been consciously attending, may have thought that the only reasonable strategy in the classification task was to check each item against any remembered exemplars. Consequently, the similarity effect would be expected to overwhelm any contribution of grammatical status to classification behavior. By presenting a greater number of exemplars, we hoped to eliminate the bias against abstraction and use of rule-based information that we felt was inherent in the Vokey and Brooks design.

A second modification involved the orienting task instructions. Both Reber and Allen (1978) and Vokey and Brooks (Experiment 1) used deliberately vague "observation" instructions, simply asking subjects to pay attention to the exemplars. Reber and Allen argued that such instructions optimized abstraction performance, because they did not require attention to any potentially distracting information, and because categorization accuracy had been highest in this condition. Vokey and Brooks, however, found that this orienting condition resulted in the greatest advantage of similarity distance over grammatical status in controlling classification judgments. In either case, the use of such instructions diminishes experimental control and allows for substantial differences in individual encoding strategies, making the reported analyses of group classification data difficult to interpret. We, therefore, presented our study phase as a consumer marketing survey on preferences for various names for home computer models, with the exemplar letter strings as "names" to be rated. Although this is a somewhat irregular manipulation in concept formation research, it is analogous to incidental learning paradigms in memory research, such as rating the pleasantness of words. The main objective, in the present case, was to constrain the observation orienting condition to a greater extent than in the previous experiments, without explicitly inducing either a memorization or a pattern-searching strategy.

Finally, Vokey and Brooks argued that Reber and Allen's (1978) data were subject to bias in that their measure, accuracy of grammatical assignments, did not take into account unconditional response frequency. To control for this, Vokey and Brooks scored their data as

proportion of total "grammatical" and "nongrammatical" responses made to each type of stimulus. However, their overall rate of "grammatical" responding was quite low, 33% in the observation condition. This may be a result of subjects' consistent aversion to choose any item for which they could not remember a similar exemplar, and may not fully reflect their sensitivity to grammatical information. The present experiment used a forced-choice design, allowing for a more sensitive assessment of the information used for grammaticality judgments while controlling for overall response bias. Because subjects had to make a "grammatical" judgment on each trial, the data scoring was consistent with the Vokey and Brooks's method: proportion of "grammatical" responses to each stimulus type.

In sum, these alterations were designed to provide a more accurate portrayal of the sources of information available for classification of new items, without biasing subjects to use either an analogy or abstraction strategy. Two additional tasks, item completion and recognition, were included to test the hypothesis that the former would reflect the ability to use abstracted rules and the latter would assess retention of individual orienting exemplars on which an analogy strategy could be based.

METHOD

Materials

Classification task. The study-phase items were 16 letter strings, varying in length from three to eight letters, chosen from the pool of grammatical strings generated from the grammar shown in Figure 1. These exemplars were selected to provide a representative sample of all possible rules and subrules, with an approximately equal number of possible initial, terminal, and repeating letters included. Each string was typed on a separate 4×6 in. card, and strings of approximately equal length were paired for presentation, resulting in 8 pairs of items. Three repetitions of the 8 pairs were constructed, with the order of pairs randomized in each set. Thus, the full set of orienting exemplars consisted of 24 pairs of grammatical letter strings, presented in a predetermined random order.

Classification test strings were constructed for 12 of the 16 exemplars. (Given the constraints of the design, appropriate test items were not available for the other 4 study strings.) As shown in Table 1 below, 4 different test strings were constructed for each of these 12 items. Two of the strings were generated by the finite-state grammar (G items) and two were formed by violating the grammatical rules (NG items). For the grammatical letter strings, one was highly similar to the exemplar, differing only by the substitution of a single letter (CL item), and one was less similar, differing in three or more letter positions (FAR item). Each FAR item differed from all 16 exemplars according to the same criteria. The nongrammatical strings were constructed in the same way, with rule violations introduced at those positions used for substitution in the grammatical items. Thus, each of the 48 test strings was one of four types, defined by the orthogonal manipulation of grammatical status (G vs. NG) and similarity distance (CL vs. FAR), distributed as shown in Table 1.

For the forced-choice classification test, 48 pairs of stimuli were formed from this pool of letter strings, with an equal representation of the four item types shown above. Thus, each grammatical item was presented twice, paired with a different type of nongrammatical string on each trial. To prevent subjects from attending to only a single-letter position, there were no pairs with two close items derived from the same exemplar. The position of an item in a pair

Table 1
Examples of Classification-Task Letter Strings

Orienting Exemplars	Test Items			
	G-CL	NG-CL	G-FAR	NG-FAR
VXRRR	VXRRM	VXRRX	VXTVX	VXTVM
MVRXM	MVRXR	MVRXT	MTVX	MTVR
(16)	(12)	(12)	(12)	(12)

Test Pairs			
G-CL/NG-CL	G-CL/NG-FAR	G-FAR/NG-CL	G-FAR/NG-FAR
VXRRM/MVRXT	MVRXR/VXTVM	VXTVX/MVRXT	MTVX/VXTVM
(12)	(12)	(12)	(12)

Note—numbers in parentheses represent actual number of each stimulus type.

was counterbalanced across presentations. Each pair was typed on an individual 4×5 in. page, and the order of the 48 trials was randomized, producing a single test order given to all subjects in booklet form.

Completion task. Five exemplars presented in the study phase of the classification task and five grammatical items which had never been presented were selected. For each string, one or two letters, distributed across all positions in the entire set, were omitted and replaced by a blank. The order of letter strings was randomized, and all 10 were typed on a single page.

Recognition task. A yes/no recognition test list was constructed, using 12 studied exemplars and 12 grammatical letter strings not previously presented. Items were typed on a single page in a predetermined random order, subject to the restriction of no more than three consecutive old or new items.

Subjects and Procedure

Thirty university undergraduates participated in the experiment for either course credit or payment. All subjects were tested individually in a single session lasting approximately 45 min. The order of tasks was identical for all subjects.

Subjects were told that the study phase of the classification task was a brief marketing survey in which they were to indicate preference for various letter strings as home-computer names. Each participant was given a rating form headed by the columns “preferred” and “rejected” and was instructed to copy the letter strings from each pair under one or the other column, depending on his/her preference. Each pair of items was presented individually, remaining in view until both letter strings had been copied onto the rating form. Upon completion of the 24 self-paced rating trials (lasting approximately 8 min), the subject was asked for comments concerning the rating process.

Next, subjects were informed about the rule-based nature of the items just rated, but not about the rules themselves. They were told that a number of new pairs of letter strings would be presented, with one member of the pair conforming to the rules which generated the exemplars and the other string violating these rules. They were to decide which item followed the same set of rules, or grammar, as the previously rated strings, and to indicate their choice by circling the grammatical item. Subjects were then given the booklet of 48 item pairs, and allowed to proceed through it at their own pace, but prevented from reviewing previous trials. Upon completion of this test phase, we solicited subjects’ introspections about the strategies they had used to judge items.

The 10-item completion task was given next, with subjects filling in each letter string according to what they believed were the correct grammatical rules. This was followed by a 5-min distraction interval of arithmetic problem-solving. The recognition test was given last; individuals were instructed to circle items which they remembered viewing during the initial rating phase of the study.

RESULTS AND DISCUSSION

An unexpectedly large degree of between-subject variability was observed in the data. Thus analyses were carried out for both group data, including all 30 subjects, and for subgroups whose determination will be discussed below. We will report first the overall results (n=30) for each task, and then discuss results and analyses for the postexperimentally determined subgroups. The former data are presented in the rows designated “Overall” in the appropriate tables.

For the classification task, the mean proportion of items labeled “grammatical” for each of the four letter-string types (G-CL, NG-CL, G-FAR, NG-FAR) are shown in Table 2A. These data were submitted to a 2 (grammatical status) × 2 (similarity distance) ANOVA. Both main effects were statistically significant [F(1,29) = 8.02, p < .01 for grammatical status and F(1,29) = 13.12, p < .002 for similarity distance]. The interaction effect was not significant. To explore these effects further, the classification data were reduced in the following manner. For each subject, a measure of the effect of grammatical status (GSE) on judgments was determined by subtracting the number of NG items labeled “grammatical” from the number of G items labeled “grammatical.” Likewise, the effect of similarity distance (SDE) was assessed by subtracting the “grammatical” responses to FAR items from “grammatical” responses to CL items. Finally, the difference between GSE and SDE for each subject provided an index of the strength of grammatical status with similarity distance partialled out (GRAM SC). As shown in Table 2B, the magnitudes of the grammatical

Table 2A
Mean Percentage of Test Items Labeled “Grammatical” as a Function of Grammatical Status and Similarity Distance

	G		NG		Percent Correct
	CL	FAR	CL	FAR	
Overall (n=30)	56.81	49.31	51.25	42.64	53.06
Grammaticality Group	56.11	58.61	45.28	40.00	57.36
Similarity Group	57.50	40.00	57.22	45.28	45.75

Table 2B
Response Difference Scores Indicating Strength of Grammatical Status, Similarity Distance, and Grammatical Status Without Similarity Distance in Controlling "Grammatical" Assignments

	GSE	SDE	GRAM SC
Overall (n=30)	2.93	3.87	-0.93
Grammaticality Group	7.07	0.67	6.40
Similarity Group	-1.20	7.07	-8.27

Note - GSE-grammatical status; SDE-similarity distance; GRAM SC-grammatical status without similarity distance.

status effect (GSE) and the similarity distance effect (SDE) were approximately equal in the overall group data.

Our findings that both similarity distance and grammatical status contributed to classification performance replicate the findings of Vokey and Brooks (in press). However, they reported a much larger effect of similarity distance versus grammatical status, whereas the contribution of the two factors was approximately equal in the present experiment. This difference is consistent with our suggestion that their manner of presenting exemplars for study effectively biased subjects to use an analogy or remembering strategy in the classification test and that the present design eliminated this bias. Although classification accuracy was rather low (53%), our transfer test was more demanding than those used in previous studies (e.g., Reber, 1976; Reber & Allen, 1978), because rule violations were never introduced in obvious positions, such as the initial letter, and no multiple violations were used. Our study phase design also may have increased the opportunity for rule abstraction so that similarity distance was not the sole basis for classification judgments. Additional support for this conclusion is given by the following analyses.

Additional data regarding subjects' apprehension of grammatical rules and memory for exemplars was provided by the completion and recognition tasks, respectively. As shown in Table 3, mean accuracy for the completion task was 42%, which is well above the chance level of 20%, and accuracy did not differ for studied versus new letter strings. Thus, despite the rather low degree of classification accuracy shown in Table 2A, subjects' apprehension of the grammar was sufficient to allow completion of a number of novel letter strings. Although memory for specific exemplars could contribute to letter-string completion performance, reliance on this source of information should result in higher completion accuracy for studied strings than for novel items. This was not observed. However, it may be that deleting letters from studied strings effectively destroyed the "oldness" of the

Table 3
Mean Accuracy for Completion and Recognition Tasks

	% Correct Completion		d'
	New Items	Old Items	
Overall (n=30)	40.00	43.33	1.676
Grammaticality Group	44.00	42.67	1.233
Similarity Group	36.00	44.00	2.119

items. The mean d' score for recognition across subjects was 1.676, which is significant at the .05 level in a one-tailed z-test, indicating better than chance memory for the original study exemplars. Although the similarity distance of the lures in the recognition test was not systematically varied, a post hoc classification of them as close or far according to the original selection criteria failed to reveal a relation between probability of false recognitions and similarity distance.

Results from these two tasks indicate that both rule-based and specific item information were available to subjects and these two types could be used selectively for appropriate tasks. That is, subjects were able to use grammaticality information in completing both old and new letter strings and were able to discriminate old from new grammatical items in the recognition task. These results are not consistent with Vokey and Brooks's (in press) argument that the same information and processes are used in classification and recognition judgments. Our subjects were able to use rule-based information for completion of both studied and new letter strings, but they were able to reject grammatical recognition lures at above-chance accuracy. The critical point here is not that exemplars are stored and accessible, because similar results for exemplar recognition have been reported in previous studies supporting an abstraction model (e.g., Homa et al., 1981; Posner & Keele, 1970). Rather, the central question is whether exemplar memory mediates classification judgments. This issue is addressed more precisely in the following analysis.

A final analysis was conducted to examine the relation between exemplar recognition and classification behavior. For each subject, "grammatical" responses to CL classification letter strings were rescored, contingent on subsequent recognition of the exemplars from which the test strings were derived. Thus, if the subject had correctly recognized the study exemplar "VXRRR," the items "VXRRM" (G-CL) and "VXRRX" (NG-CL), both of which appeared twice in the classification test, were examined to see whether they had been labeled "grammatical" during the test. Of course, the critical comparison is for NG-CL items which were either judged grammatical or rejected, because the G-CL items could have been selected either on the basis of similarity or grammatical status. We would expect that exclusive use of an analogy strategy in classification would be reflected in an equally large labeled-rejected difference for both NG-CL and G-CL items, because both types are highly similar to remembered exemplars. Mean scores for the difference between the number of CL items labeled "grammatical" and those rejected are displayed in Table 4. The overall labeled-rejected difference was significant according to a paired-difference t-test [$t(29) = 5.21, p < .0001$], as was the labeled-rejected difference for G-CL items [$t(29) = 5.27, p < .0001$]. Across all 30 subjects, the difference for NG-CL letter strings was not significant. Thus, grammaticality information contributed substantially to performance on transfer items close to specific remembered ex-

Table 4
Mean Labeled-Rejected Difference for Recognition
Contingent CL Classification Items as a
Function of Grammatical Status

	G	NG
Overall (n=30)	3.33	1.50
Grammaticality Group	3.07	-0.73
Similarity Group	3.60	3.73

emplars. This result runs counter to the predictions of a strict classification-by-analogy model.

A surprising finding of the present experiment was the large degree of between-subject variability in the effects of grammatical status and similarity distance on classification judgments. To explore this variability further and to examine whether it reflected clear strategy differences, two subgroups of 15 subjects each were created postexperimentally. This was done by performing a median split on the group classification GRAM SCs, which assess the strength of grammatical status with similarity distance taken out. For convenience, the subgroup with the higher scores will be referred to as the Grammaticality Group, and that with the lower as the Similarity Group. Subgroup means for the classification measures are presented in Tables 2A and 2B. Mean percent correct classification assignments differs, of course, between the two subgroups [$t(28) = 5.84, p < .0001$].

The two subgroups did show differences in recognition and completion task performance, as shown in Table 3. However, only one of these between-group differences was statistically significant [for recognition d' , $t(28) = 1.796, p < .05$], with the Similarity Group making more correct recognitions. This difference in exemplar memory between the subgroups cannot account for the observed differences in similarity effects on classification judgments, however, as correlations between the similarity distance effect (SDE) in classification and recognition d' were not significant. Indeed, some subjects with very high d' scores also showed strong grammaticality effects in classification.

The clearest difference between the subgroups was found in the recognition-contingent analysis of CL classification items, as indicated by the group means shown in Table 4. Recall that the critical labeled-rejected difference is the one for NG-CL items, because the effect of grammaticality may be contributing to this difference for the G-CL letter strings. The difference in NG items was significant only for the Similarity Group [$t(14) = 3.89, p < .002$], and not for the Grammaticality Group [$t(14) = 0.51$]. This performance difference in recognition-contingent assignments for NG-CL items between the Similarity and Grammaticality Groups was significant [$t(28) = 3.26, p < .003$]. Thus, memory for a highly similar exemplar was more likely to be reflected in "grammatical" responses to NG items for only a subgroup of subjects. For these subjects, it may be the case that similar processes underlie exemplar recognition and grammaticality judgments.

In sum, it appears that for some individuals, such as those in the Similarity Group, classification judgments are based on analogy to stored exemplars. However, the present data also indicate that classification based on abstracted rules is a viable explanation of performance in the artificial-grammar learning paradigm. Thus, subjects in the Grammaticality Group were able to discriminate between grammatical and nongrammatical items that were highly similar to actual remembered exemplars. These subjects appeared to be using a qualitatively different source of information or strategy to classify the transfer items. Given these results, one is not compelled to accept Vokey and Brooks's (in press) argument that the possible confounding of similarity distance and grammatical status in previous experiments renders implausible an abstraction interpretation of categorization performance. In addition, the present data are consistent with Reber's argument that a nonconscious rule-induction process is carried out during exemplar encoding, given that our study phase instructions did not explicitly emphasize rule induction and the exemplars were not presented in a manner that made the patterns of invariance highly salient. However, the presumed strategy differences we found place limitations on a clearer description of the process and the resulting mnemonic representation, both in terms of the present data and of the results of previous experiments.

Unfortunately, the source of these strategy differences cannot be specified given the data reported here, and the subjects' protocols were not informative regarding their decision-making processes. All subjects reported that their classification judgments were guided by both overall goodness-of-fit and similarity to rated items. Subjects' locus of control has been found to interact with classification performance (Kassin & Reber, 1979), but no such extraexperimental factors were explored in the present study. A reasonable interpretation of the present data is that the changes in the study phase sufficiently constrained encoding strategies, but not classification strategies. This is supported by the fact that recognition memory for orienting exemplars was above chance across all 30 subjects, but only a subset of the group consistently used this information in assigning "grammatical" labels to test items. The reported differences in subjects' response patterns under identical study and test conditions point to a need for further efforts directed at analyzing categorization behavior in terms of both subject and task-related constraints.

In summary, the present data provide support for the hypothesis that rule-based information can be abstracted and used for classification judgments even in situations in which exemplar-based information is available to the subject as an alternate basis for responding. Thus, classification by analogy does not carry the full explanatory weight in this paradigm. The artificial grammar we used is characterized by a list of specific transformation rules, and thus the set of grammatical strings do not formally represent an ill-defined category. However, the rules

governing grammatical formation of strings were not readily available to our subjects. We base this claim on the facts that the experimental design precluded any easy apprehension of the rules (i.e., incidental learning, no blocking of strings by subrules) and that our subjects did not behave like hypothesis testers as in concept-identification studies. Because the structure and rules were not transparent to the subject, and because the data bear strong similarities to findings with ill-defined categories, we believe the present study has relevance to this broader domain of categorization.

REFERENCES

- ALLEN, R., & REBER, A. S. (1980). Very long term memory for tacit knowledge. *Cognition*, **8**, 175-185.
- BROOKS, L. R. (1978). Non-analytic concept formation and memory for instances. In E. Rosch & B. Lloyd (Eds.), *Cognition and concepts*. Hillsdale, NJ: Erlbaum.
- ELIO, R., & ANDERSON, J. R. (1981). The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning & Memory*, **7**, 397-417.
- HAYES-ROTH, B., & HAYES-ROTH, F. (1977). Concept learning and the recognition and classification of exemplars. *Journal of Verbal Learning & Verbal Behavior*, **16**, 321-338.
- HINTZMAN, D. L., & LUDLUM, G. (1980). Differential forgetting of prototypes and old instances: Simulation by an exemplar-based classification model. *Memory & Cognition*, **8**, 378-382.
- HOMA, D., STERLING, S., & TREPPEL, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning & Memory*, **7**, 418-439.
- HOMA, D., & VOSBURGH, R. (1976). Category breadth and the abstraction of prototypical information. *Journal of Experimental Psychology: Human Learning & Memory*, **2**, 322-330.
- KASSIN, S. M., & REBER, A. S. (1979). Locus of control and the learning of an artificial language. *Journal of Research in Personality*, **13**, 112-118.
- MEDIN, D. L., & SCHAFFER, M. M. (1978). Context theory of classification learning. *Psychological Review*, **85**, 207-238.
- MEDIN, D. L., & SMITH, E. E. (1981). Strategies and classification learning. *Journal of Experimental Psychology: Human Learning & Memory*, **7**, 241-253.
- POSNER, M. I., & KEELE, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, **77**, 353-363.
- POSNER, M. I., & KEELE, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, **83**, 304-308.
- 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 artificial grammars: The role of instructional set. *Journal of Experimental Psychology: Human Learning & Memory*, **2**, 88-94.
- REBER, A. S., & ALLEN, R. (1978). Analogy and abstraction strategies in synthetic grammar learning: A functional interpretation. *Cognition*, **6**, 189-221.
- REBER, A. S., ALLEN, R., & REGAN, S. (1985). Syntactical learning and judgment, still unconscious and still abstract: Comment on Dulany, Carlson & Dewey. *Journal of Experimental Psychology: General*, **114**, 17-24.
- REBER, A. S., KASSIN, S. M., LEWIS, S., & CANTOR, G. (1980). On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 492-502.
- REED, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, **3**, 382-407.
- STRANGE, W., KENNEY, T., KESSEL, F. S., & JENKINS, J. J. (1970). Abstraction over time of prototypes from distortions of random dot patterns: A replication. *Journal of Experimental Psychology*, **83**, 508-510.
- VOKEY, J. R., & BROOKS, L. R. (in press). Taming the clever unconscious: Analogic and abstraction strategies in artificial grammar learning. *Cognition*.

(Manuscript received December 5, 1984;
revision accepted for publication May 8, 1985.)