Journal of Experimental Psychology:
Learning, Memory, and Cognition
1995, Vol. 21, No. 4, 899-912

Copyright 1995 by the American Psychological Association, Inc.
0278-7393/95/$3.00

Modality Independence of Implicitly Learned Grammatical Knowledge

Gerry T. M. Altmann, Zoltán Dienes, and Alastair Goode
University of Sussex

In each of 4 experiments exemplars of an artificial grammar were presented in one modality (e.g., sequences of tones differing in pitch in Experiments 1 and 2, or sequences of spoken syllables in Experiment 3, or sequences of arbitrary graphical symbols in Experiment 4), but the subsequent classification task was performed on novel sequences in another (e.g., sequences of letters in Experiments 1 and 2, sequences of arbitrary graphic symbols in Experiment 3, and sequences of written syllables in Experiment 4). Prior exposure to the grammar improved classification of novel stimuli, even across modalities. Participants who received either no preexposure or were exposed to pseudorandom sequences showed no such improvement. Consequently, part of the learning process can take place prior to any exposure to the domain within which categorization is to take place.

Knowledge of the regularities underlying variation in the external environment plays a central role in the workings of the human cognitive system. Without this knowledge, the cognitive system would be unable to perform even the simplest of tasks. The acquisition of such knowledge is thus crucial to the emergence of cognitive ability. The aim of this article is to address, within the context of artificial grammar learning, a specific issue concerning the acquisition process: Is acquisition modality specific, or can knowledge of regular variation in one modality underlie sensitivity to regular variation in another? Thus we ask whether modality dependence (and indeed domain dependence, defined below) constrains the acquisition and application of the grammatical knowledge.

In a number of studies, Reber and others have demonstrated that exposure to sequences of letters generated by an artificial grammar enables participants to perform at significantly above chance when subsequently discriminating between new sequences that either do or do not obey the rules of the grammar (e.g., Dienes, Broadbent, & Berry, 1991; Mathews et al., 1989; Reber, 1967, 1969). Despite performing above chance, participants were generally unable to freely report the rules they had used during this classification task, and on those occasions when they could, the rules they described did not adequately explain their improved performance (Dienes et al., 1991; Mathews et al., 1989). Participants who either had not received prior exposure to any grammatical exemplars (cf. Mathews et al., 1989) or were given exposure to sequences generated by a different grammar (cf. Dienes et al., 1991; Reber, 1969) did not perform as well when given the same discrimination task. These studies demonstrated that knowledge about some subset of the variation permitted by an artificial grammar can be learned, without the need to expose participants to the actual rules of that grammar. All that is required is to expose participants to the consequences of those rules. In the remainder of this article we refer to the knowledge that has been acquired as *the grammar,* and in the context of grammar learning this expression is thus shorthand for *knowledge about some subset of the lawful variation in the input.*

Recently, there has been considerable interest in elucidating the nature of the knowledge that is abstracted from the exemplars given in the preexposure phase. Reber (1969) argued that what is acquired is abstract grammatical knowledge; Perruchet and Pacteau (1990) have suggested instead that participants are sensitive to little more than knowledge about the particular bigrams that can occur in the letter sequences typical of these studies; Mathews (1990) argued that the knowledge acquired is in fact far richer than simple bigrams and includes information concerning the spatial positioning of these bigrams relative to other elements in the strings (i.e., the knowledge is perhaps better classified in terms of distributions of "ngrams"). The representational issue is complicated by the fact that whatever the nature of the knowledge that is acquired in artificial grammar learning, it is remarkably insensitive to changes in the "vocabulary" on which basis the knowledge was first acquired. Reber (1969) and Mathews et al. (1989) have shown that so long as the

---

Gerry T. M. Altmann, Zoltán Dienes, and Alastair Goode, Laboratory of Experimental Psychology, University of Sussex, Brighton, United Kingdom. The work reported in this article was supported by Joint Council Initiative Grant G9110951 and by the United Kingdom component of the Human Frontiers in Science Program project on Processing Consequences of Language Phonologies. Aspects of the work were formulated while Gerry T. M. Altmann held a University of Melbourne Visiting Fellowship.

Preliminary data were first presented at The Trieste Encounters on Cognitive Sciences, June 1991; the Seventh Australian Language and Speech Conference, November 1991; and the 21st Annual University of Wisconsin-Milwaukee Linguistics Symposium on the Reality of Linguistic Rules, April 1992.

We are grateful to Astrid Shepman for her help in devising the auditory stimuli for Experiment 3 and to Helena Davidson for her help in devising the ungrammatical stimuli used in Experiment 4. Finally, we are grateful also for the constructive comments on a version of this article that we received from Axel Buchner and Robert Mathews. Any errors of interpretation are our own.

Since writing this article, we have learned that Louis Manza and Arthur Reber, working independently and without any knowledge of our own studies, have found effects similar to the ones we report here (Manza & Reber, 1991, 1994).

Correspondence concerning this article should be addressed to Gerry T. M. Altmann, Laboratory of Experimental Psychology, University of Sussex, Brighton BN1 9QG United Kingdom. Electronic mail may be sent via Internet to gerrya@epunix.sussex.ac.uk.

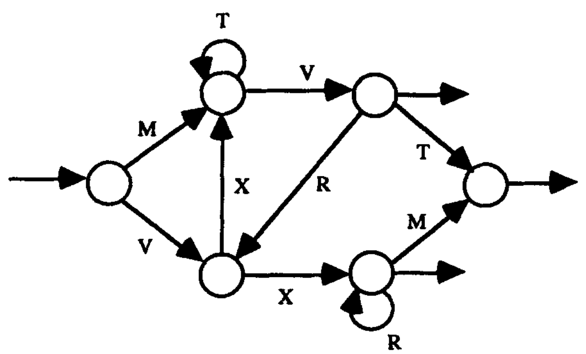

Figure 1. The finite-state grammar used in Experiments 1 and 2.

underlying grammars are the same, it matters little whether the letters used to make up the exemplars in the learning phase are the same as those used as the basis for establishing that learning has taken place (whether in a subsequent classification task—Mathews et al.—or in a continuation of a memorization task, with errors to criterial learning as the dependent variable—Reber). In other words, changing the vocabulary makes little difference if the underlying syntax remains the same. Clearly, a model that is based on the extraction of the distributional characteristics of the individual, specific, letters would not be enough to explain this transfer effect. The acquired knowledge might be better characterized in terms of the distributional characteristics of something akin to form classes (in which the individual letters are viewed as words belonging to a syntactic form class), or it might be better characterized in terms of an abstract grammar (with the same assumption concerning the relationship between letters and syntactic form class). In either case, one might argue that some level of underlying structure is implicated. On the other hand, Brooks and Vokey (1991) suggested that transfer between letter strings from different vocabularies could have taken place in these studies partly because of some abstract similarity between the letter strings in the test phase and stored representations of the individual exemplars presented in the learning phase, with transfer across letter sets being due to the formation of abstract analogies computed separately on an item-by-item basis (cf. Whittlesea & Dorken, 1993). Before returning to this issue (we present data inconsistent with a strict item-by-item account in Experiment 3), we consider an extension of the original Reber finding.

The finding that vocabulary sets can be changed between the learning and classification phases of these studies raises an interesting and important question: What relationship between the vocabulary sets must obtain for learning to transfer between them, other than a common underlying structure (a requirement even if transfer is by analogy)? Another way of framing this same question is to ask whether, and to what extent, the knowledge that is acquired in implicit artificial grammar learning is available for use across different domains. For present purposes, we can define two input domains as being different to the extent that the vocabularies used to describe the input signals do themselves differ, and (crucially) *the relevant mapping between them is not known a priori.* Thus, if two input signals can only be described by using different sets of descriptors, and there is no a priori relationship between the elements of the different sets, then we define the two signals as belonging to different domains. A strong example of cross-domain transfer of implicit knowledge would be one in which the two domains occur in different modalities. For instance, can knowledge about which patterns are permissible in the visual (e.g., orthographic) domain influence the process of discerning patterns in the auditory (e.g., tonal) domain? That is, will exposure to a set of grammatical exemplars in one modality improve classification of a test set (relative to controls with either no prior exposure or exposure to a different grammar) in the other modality?

Experiment 1

In this first experiment, a grammar generated simple melodies or sequences of letters. The aim was to see whether prior exposure to, for instance, the melodic exemplars would improve classification of the letter sequences, and vice versa. Classification consisted of distinguishing between sequences that were either grammatical or ungrammatical. Two control conditions were used: one in which participants received no prior exposure, and another (used only in the case of prior exposure to melodies and subsequent classification of letter sequences) in which participants were preexposed to exemplars from a different grammar. If cross-domain transfer takes place, then the participants given prior exposure to the grammar, irrespective of the modality in which this takes place, should show improved classification relative to participants in the two control conditions.

The precise translation used to map melodic sequences onto letter sequences is of critical importance: If tonal distance mapped onto alphabetic distance (e.g., the letter sequence *ABC* mapping onto the melodic sequence cde, *BCD* onto def, and so on), the mapping between tones and letters could be sufficiently transparent that the sequences would violate the criterion for constituting separate domains (described above). We therefore ensured that the mapping between tones and letters was random and that alphabetic distance did not, and in fact could not, correspond to tonal distance.

Method

Participants. Forty-six University of Sussex undergraduates participated in the study.

Stimuli. The grammar was identical to that used in a number of earlier implicit learning studies (e.g., Dienes et al., 1991; Dulany, Carlson, and Dewey, 1984; Reber, 1967) and is shown in Figure 1. The same stimuli were used as given in Dienes et al. and Dulany et al.

The grammar shown in Figure 1 was used to generate 40 sequences that varied in length from between three and six letters. Twenty of these were assigned to the learning set, and 20 to the classification set. Five of the 20 exemplars in the learning set were added to the classification test set, making a total of 25 test stimuli. A further 25 sequences were created that were ungrammatical and were matched against the grammatical sequences for length and frequency of occurrence of individual letters. These 25 ungrammatical sequences were randomly interspersed amongst the 25 grammatical sequences (also in random order) to create the full classification set. The two sets of letter sequences were then translated into musical tones by using a random mapping of the tones to letters (c-*M*, d-*T*, e-*R*, g-*V*, a-*X*, in which lowercased letters indicate the musical note, with c referring to Middle C). We generated the melodies using a sampling keyboard (generating simple sine waves at the appropriate frequencies at approximately two beats per second). A final set of 20 ungrammatical musical sequences, in which successive notes were chosen at random, was also generated. These sequences were matched against the others for length and frequency of occurrence of individual notes. We used this set as a control condition in which participants are exposed to exemplars from a grammar (or more correctly, a set of grammars each capable of generating these exemplars) that differs from that which was used to generate the target sequences in the classification set. We used this condition as an additional control against which to compare participants who were given exposure in the auditory domain before testing in the orthographic domain.

Design. We used a mixed design with four groups. One group ($n = 12$) acted as a control and was given no learning phase and proceeded directly to the classification task. Two further groups ($n = 12$ each) were exposed in the learning phase to either the letters or the tones that had been generated by the same grammar. The same (random) order of sequences was used for both the melodies and the letter sequences and was fixed across all participants. All three groups were given both classification tasks (i.e., tones and letters). Half of the participants in each group were given the tones to classify first, and half were given the letters to classify first. Thus, order of classification task was fully counterbalanced. Subsequent analyses revealed that there were no main effects of test order nor any interactions with test order (all $ps > 0.1$); we therefore omit further discussion of this counterbalancing procedure. To minimize any order effects in the classification stimuli, we had half of the participants classify the stimuli within each modality in one order, and we had the other half classify the stimuli in the reverse order. No participant was given the same ordering across the two modalities. A final group of participants ($n = 10$) was exposed in the learning phase to the random tones and was then required to classify only the orthographic sequences. This condition was intended as a further control in the event that transfer occurred between tones and letters. Again, half of the participants classified the sequences in one order, and the other half classified the sequences in the other order.

Procedure. Participants were presented with the musical stimuli over headphones, with approximately 5 s between each stimulus sequence. The tape ran for approximately 3 min. The letter sequences were presented on a single sheet of paper. Except for the group that only took part in the classification task, and depending on which of the remaining three groups they were in, participants were first asked either to listen to 20 melodies or to study the 20 letter sequences appearing on a single sheet of paper. Participants were allowed to inspect the letter sequences for between 3 and 4 min. If they inquired as to the purpose of this, they were informed that they would be asked some questions about the stimuli later. After they had listened to the tape, or inspected the training set, they were told that the 20 exemplars they had seen, or heard, had been generated by a computer program. They were then told that they would now be given two sets of 50 sequences each and that half of these sequences were generated by the same program and half were generated at random. Their task was to classify which sequences had been generated by the same program and which were random. When participants had to classify the melodic stimuli, they wrote down their response after hearing each sequence. When participants had to classify the letter sequences, they wrote their response against each of the 50 sequences appearing on the response sheet. Participants in the no-learning conditions were given both the 50 auditory stimuli and the 50 letter sequences in counterbalanced order. They were told that half of the sequences had been generated by the same program and that half had been generated at random. Their task was to classify which were which.

Table 1
Percentage Correct of Classification Scores as a Function of Learning Set and Test Set

	Test set			
	Letters		Tones	
Learning set	%	*SD*	%	*SD*
Control (no learning)	50	1.1	49	1.0
No transfer (same as test)	59	1.7	57	1.2
Transfer (different to test)	54	1.3	56	0.8
Random tones	48	1.2		

Results

The percentage correct of classification scores are shown in Table 1.

A two-way analysis of variance (ANOVA) with variables test set (letters and tones) and learning set (none, same modality as test, and different modality from test) was performed on all but the random tone data. There was no effect of test set ($F < 1.0$), an effect of learning set, $F(2, 33) = 26.20, p = .0001$, $MSE = 19.60$, and no interaction between learning and test set, $F(2, 33) = 1.67, p > .2, MSE = 16.50$. Planned comparisons revealed significantly improved classification in the no-transfer conditions relative to the no-learning controls, $F(1, 33) = 51.44, p = .0001, MSE = 19.60$, from an average of 49% to an average of 58%. Comparisons also revealed significantly improved classification after transfer relative to the no-learning controls, $F(1, 33) = 19.66, p = .0001, MSE = 19.60$, from an average of 49% to an average of 55%. The comparison between the transfer (55%) and no-transfer (58%) conditions was also significant, $F(1, 33) = 7.50, p < .01, MSE = 19.60$. Thus the no-transfer conditions (letters to letters and tones to tones) lead to better classification in the test phase than the transfer conditions (letters to tones and tones to letters). Not surprisingly, given that there was no main effect or interaction involving test set in the original ANOVA, no planned comparisons interacted with test set (all $ps > .1$). Finally, inspection of Table 1 shows that classification of letters when preceded by exposure to random tones was nonsignificantly lower than when preceded by no exposure at all (48% vs. 50%).

A further analysis that would be useful would be one that assessed whether there were any trends in correct classification of a transfer test item against serial position of the test item. If participants build up a mapping over the course of the test phase, classification performance should gradually improve. If participants classify each test stimulus separately according to some abstract (conscious or unconscious) analogy (cf. Brooks & Vokey, 1991), there should be no trend. However, with overall performance in the transfer conditions at only 54.5%, we simply lack the power to distinguish the alternatives, and, in fact, there was no observable trend (b = .003). A better test, and one that is currently under way, is to compare performance in this experiment with one in which the mapping between tone and letter is changed for each sequence in the test phase. If some form of correspondence is established separately for each sequence, then changing the mapping should not affect performance. If instead the system attempts

Table 2
Percentage Correct of Classification Scores as a Function of Learning Set and Test Set

	Test set			
	Letters		Tones	
Learning set	%	*SD*	%	*SD*
Control	50	1.1	50	1.9
Transfer	54	2.0	56	2.0

to establish an internal mapping that is consistent across stimuli, then changing the actual mapping should affect performance. Whittlesea and Dorken (1993) independently reported testing a group with randomly changed mappings, but it is not possible to say whether they found transfer as they did not use a control group against which to compare performance.

Discussion

We designed Experiment 1 to test whether the correct classification of sequences in one modality would be facilitated by prior exposure to grammatical exemplars presented in another modality. We found, relative to the control conditions, that classification was significantly improved, even when a switch in modality had occurred (the transfer conditions). Moreover, transfer occurred irrespective of test modality and without feedback during the test phase. It would appear, therefore, that cross-domain transfer of acquired grammatical structure does occur.

The inclusion of the condition in which random melodies were used demonstrates that the "content" of the melodic exemplars was crucial to establishing a transfer effect. Indeed, this comparison is directly equivalent to Reber's (1969) demonstration that transfer can occur across changes in vocabulary but not across changes in the underlying syntax. We note that Perruchet (1994) argued that many previous demonstrations of transfer effects have failed to use any control group against which to compare performance in the transfer condition(s). In the present study, we, in fact, have two such groups (and we note also that Brooks and Vokey's, 1991, demonstration of transfer included a counterbalancing procedure that effectively provided a control comparison).

Overall, the effects of learning were small, even when no transfer was involved: Classification in the no-transfer conditions increased by an average of only 9%, although this is comparable with other studies (e.g., Dulany et al., 1984) that showed similar increases of around 10% from baseline. The effects of learning, and transfer, were nonetheless significant, albeit small. Although the transfer conditions led to an absolute increase in performance, relative to the controls, of only 6% (i.e., from an average of 49% in the control conditions to an average of 55% in the transfer conditions), the magnitude of this effect should be viewed relative to the improvement in performance seen in the no-transfer conditions (10%). In these terms, the improvement seen in the transfer conditions represents 62% of the improvement possible in the absence of transfer. Also, in the best case, from letters to tones, the corresponding figure is 83% (i.e., 7% for letters to tones relative to 9% for tones to tones). In Experiments 3 and 4, we report even larger levels of transfer.

The evidence supports the supposition, then, that grammatical structure acquired in one modality can influence the recognition of grammatical structure in another. Before considering further the possible mechanisms of such transfer, and the implications of a processor capable of such a feat, we describe some further experiments aimed at replicating and extending the results of Experiment 1 and demonstrating that a wide range of auditory or visual stimuli give rise to the same effects.

Experiment 2

The aim of Experiment 2 was to replicate the pattern found in Experiment 1, but by using a subjectively quite different auditory stimulus. In Experiment 1, participants were exposed to sequences of discrete items in either the auditory or orthographic domains. The orthographic domain is, in many languages, normally characterized in terms of such sequences, but the auditory domain need not be so constrained (and is rarely discrete). This, therefore, permits the possibility of exploring whether the human processor can identify the patterns in a continuously varying signal and use those patterns as the basis for categorizing sequences of discrete elements in another modality.

Method

Participants. Forty-eight University of Sussex undergraduates participated in the study.

Stimuli. The same letter stimuli were used as in Experiment 1. To generate the continuously varying melodic sequences, we modified the tones from Experiment 1 by changing the portamento on the synthesizer so that each note "blended" into the next. The portamento was set relative to the beat (two per second, as in Experiment 1) such that each blend would just achieve the target pitch before moving off toward the next note. Subjectively, these stimuli did not correspond to sequences of discrete pitch peaks, and there were no physical discontinuities in the transitions between the actual peaks. Thus, for successful transfer to occur it would be necessary for participants to "segment" these melodic sequences (i.e., to perceptually encode them) into sequences of discrete peaks.

Design. We used a between-subjects design with 12 participants in each of the four groups. One group acted as a control and simply classified the melodic sequences without any prior learning phase. We used the data from Experiment 1 to provide the equivalent control for the letter sequences. Two further groups were given either the letter sequences in the learning phase and the melodic sequences in the test phase or the melodic sequences in the learning phase and the letter sequences in the test phase.

Procedure. The same procedure was used as in Experiment 1.

Results

The percentage correct of classification scores are shown in Table 2.

A two-way ANOVA with variables test type (letters and tones) and learning (no learning and transfer) was performed on the data. There was no main effect of test type ($F < 1.0$) on classification scores and no interaction between test type and learning ($F < 1.0$). There was, however, a significant main effect of learning, $F(1, 44) = 8.18, p < .007, MSE = 36.70$.

Discussion

The data show not only that learning transfers from continuous domains to discrete domains (and vice versa) but also, by comparison with Experiment 1, that there was no detrimental effect of transforming the sequences of discrete tones into a signal with continuously varying pitch. The regularities underlying the two kinds of signal were equally well applied to the task of classifying letter sequences. Similarly, comparison across the two experiments demonstrates that the classification of the two kinds of signal, after prior exposure to the letter sequences, was no different. In other words, the transfer of grammatical knowledge occurs, irrespective of a change in modality and irrespective of whether the stimuli differ, across the modalities, with respect to being composed of discrete or continuously varying signals.

In Experiment 3 we describe an experiment in which a grammar was used to generate sequences of spoken consonant-vowel-consonant syllables. The grammar was a small phrase structure grammar (from Morgan, Meier, & Newport, 1987), and unlike in Experiments 1 and 2, we assigned more than one token to some of the grammatical categories generated by the grammar. After listening to sequences of syllables, participants were asked to classify sequences of arbitrary graphic symbols. As we pointed out in our introduction to Experiment 1, letters and tones can each be mapped onto ordinal scales, and there thus exists at least one mapping that is potentially transparent. Although we ensured that the actual mapping used in Experiments 1 and 2 violated ordinal structure, the stimuli used in Experiment 3 eliminate any possibility that there might exist any transparent, or a priori, mapping between the tokens of the learning domain (spoken syllables) and the tokens of the test domain (graphic symbols).

Experiment 3

Morgan et al. (1987) described a study (their Experiment 1) in which participants were presented with spoken sentences consisting of monosyllabic nonsense words generated by an artificial phrase structure grammar. Participants were simultaneously presented with a printed version of the sentence and with a sequence of reference figures (i.e., some graphic symbol to which each nonsense word referred in the artificial world generated by the grammar). Participants were given explicit instructions to "discover how the words and the figures in this world were paired and to search for patterns in the arrangements of words" (Morgan et al., 1987, p. 512). Morgan et al. found that participants could all perform virtually perfectly in a subsequent vocabulary test (matching syllables to their graphic referents) and could also perform above chance in a subsequent grammaticality test. In Experiment 3, we conducted a modified version of this task: Participants were not given any explicit instruction to search for patterns; the auditory sequences were not paired with any printed version of either the syllables themselves or their graphic referents; and in the test phase, participants were given only sequences of the graphic referents to classify as either grammatical or ungrammatical.

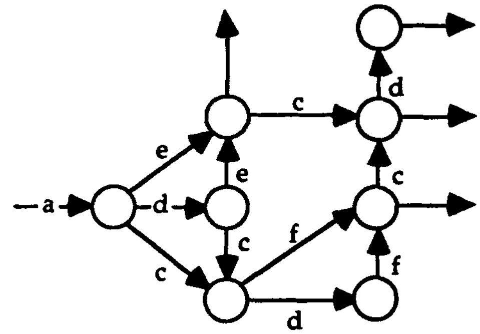

Figure 2. The finite-state grammar used in Experiments 3 and 4.

Method

Participants. Twenty-four University of Sussex undergraduates participated in the study.

Stimuli. Seventy sequences of syllables were generated from the following phrase structure grammar (from Morgan et al., 1987), a finite state version of which is shown in Figure 2.

S→ A B (C)
A→ a (d)
B→ C f
B→ e
C→ c(d)

Monosyllabic nonsense words were assigned to the terminal categories as follows (some of these syllables differ from those used in Morgan et al.):

a→ {hes or vot}
c→ pel
d→ jix
e→ {rud or sog}
f→ {kav or dup}

Thus, the sequence of Categories A D E C would translate into any of the following: *hes jix rud pel, vot jix rud pel, hes jix sog pel,* or *vot jix sog pel.*

Thirty of the 70 sequences were assigned to the learning set. The frequency of occurrence of individual syllables, and the frequency of occurrence of sequences of different lengths, were kept constant, proportionally, across the learning and test sets. Spoken versions of the 30 sequences were recorded onto DAT tape (by Gerry T. M. Altmann). Approximately 2.5 s separated each sequence, and the sequences were uttered at approximately two syllables per second. Each sequence was uttered with an intonation appropriate to its constituent structure (cf. Condition 3 of Morgan et al.'s, 1987, Experiment 1, "prosody consistent with phrase structure").

The 40 grammatical test sequences were matched with 40 ungrammatical test sequences. These were created so that, overall, they shared the length and first-order frequency statistics of the grammatical sequences (discussed below as well). The 80 test sequences were then translated into sequences of graphic symbols according to the mapping shown in Figure 3. The symbols differ from those used by Morgan et al. (1987). The assignment of symbols to syllables was random as was the final order of presentation of the sequences.

Participants. Twenty-four University of Sussex undergraduates participated in the study.

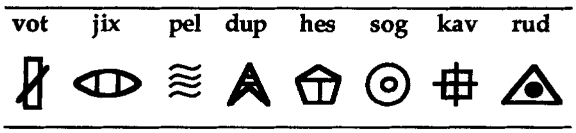

Figure 3. The mapping of nonsense words to symbols used in Experiment 3.

Design. We used a between-subjects design with 12 participants in each of two groups. One group acted as a control and simply classified the symbol sequences without any prior learning phase. The experimental group was first presented with the auditory stimuli and was then given the visual stimuli to classify.

Procedure. The experimental group was presented in the learning phase with four blocks of auditory stimuli, each block consisting of the same 30 stimuli (the learning set) presented in a different random order (although order was constant across participants). Stimuli were presented over headphones, and presentation of the four blocks of 30 stimuli took approximately 10 min. In the test phase (taken by both the experimental and control groups), participants were given response sheets on which the 80 symbol sequences appeared. To counterbalance for possible fatigue effects, we split the 80 sequences into two blocks of 40; half of the participants were given the two blocks in one order, and the other half were given the two blocks in the other order (e.g., *AB* vs. *BA*). Subsequent analyses revealed no effect of the order of the blocks on performance; we therefore omit further discussion of this manipulation. Participants' instructions were identical to those given in Experiments 1 and 2.

Results

The percentage correct of classification scores are shown in Table 3.

A one-way ANOVA (control vs. transfer) confirmed that the difference of 11% between the control and transfer group was significant, $F(1, 22) = 7.83$, $MSE = 97.00$, $p = .01$. To determine whether the transfer effect was due to just a few items, we calculated, for each test string, the number of participants giving a correct response in the transfer condition and the control condition. A difference between these conditions of 4 participants (given 12 per group) is just significant at the .05 level by using the normal approximation to the binomial, and we could expect four such items ($.05 \cdot 80$) through chance alone. In fact, 21 of the 80 test strings showed significant transfer according to this criterion. That is, the transfer effect did not appear to be due to only a few items; a substantial number of items individually showed transfer. In addition, we had asked each participant at the end of the experiment for a verbal report as to their strategies for choosing between grammatical and ungrammatical exemplars. Participants were unable to give any indication as to the basis of their decisions, or the mapping between syllables and symbols, in free report.

Discussion

Experiment 3 is significant because it demonstrates effects of transfer between sets of stimuli between which there can exist no a priori mapping. In this respect, this constitutes the *strongest demonstration* thus far of cross-domain transfer of acquired knowledge.

The data from Experiment 3 have interesting implications for the interpretation of aspects of the Morgan et al. (1987) results. They demonstrate that knowledge of the permissible sequences of syllables and/or symbols (the referents of the vocabulary items used in Morgan et al.) can be derived without the need for explicit instruction (cf. Morgan et al.) to look for appropriate patterns. This is, of course, unsurprising given the previous studies on implicit artificial grammar learning. However, the results of Experiment 3 also demonstrate that the acquisition of knowledge of the permissible orderings of the graphic referents (the symbols) does not require the simultaneous pairing of sequences of syllables with sequences of their graphic referents—this pairing is unnecessary for a mapping between the two sets of stimuli to be established. Finally, Experiment 3 demonstrates that at least some proportion of the above-chance classification of the written versions of the spoken syllables requires no a priori knowledge of the mapping between the written and auditory versions. Perhaps this point needs further spelling out: In the Morgan et al. study, participants were presented with sequences of spoken syllables in the learning phase and with sequences of written versions of these syllables in the test phase. One might then ask whether the ability to perform above chance in the Morgan et al. study relied on knowledge of the spelling rules that map spoken syllables onto their written forms. The results from Experiment 3 demonstrate that above-chance performance in such tasks requires no a priori knowledge of spelling whatsoever.

Experiment 4

Experiment 4 serves as a final replication of the transfer effects found in the preceding experiments. We used the same grammatical stimuli as used in Experiment 3, but this time the learning stimuli were sequences of symbols (corresponding to the same learning set used in Experiment 3), and we used a new set of ungrammatical stimuli. In addition, and unlike in Experiment 3, participants were tested both on sequences of (written) syllables and on sequences of symbols, allowing replication of the finding that a penalty in classification performance is incurred if the test domain is different from the learning domain.

Method

Participants. Twenty-four University of Sussex undergraduates participated in the study.

Stimuli. The 30 learning and 40 (grammatical) test sequences were the same as those used in Experiment 3. The 30 learning sequences were translated into symbol sequences according to the mapping shown in Figure 3. For each grammatical test sequence, an ungrammatical sequence was generated by reordering the elements in the

Table 3
Percentage Correct of Classification Scores as a Function of Participant Group

Participant group	%	*SD*
Control	47	1.77
Transfer	58	2.23

sequence. This guaranteed preserving the exact frequency distribution of the individual elements, the length of the individual sequences, and the frequency-by-length properties of the grammatical and ungrammatical test items. The resulting 80 test sequences were then translated into sequences of graphic symbols according to the mapping shown in Figure 3.

Design. We used a between-subjects design with 12 participants in each of the two groups. One group acted as a control and classified both the syllable and symbol test sequences without any prior learning phase. Half of the participants were given the 80 symbol sequences to classify first, and half were given the 80 syllable sequences to classify first. The order of the sequences within each test set differed (but was constant across participants) and differed from the order of presentation used in Experiment 3. The second, experimental, group was first presented with the training (symbol) sequences. Like the control group, this group also classified both the syllable and the symbol sequences, with order of classification task fully counterbalanced. Subsequent analyses revealed that there were no main effects of test order nor any interactions with test order (all $ps > 0.1$); we therefore omit further discussion of this counterbalancing procedure.

Procedure. The experimental group was presented in the learning phase with four blocks of the same 30 symbol sequences (the order of presentation of the ensuing 120 stimuli was the same as that used in Experiment 3). The stimuli were presented on seven sheets of paper, with approximately 18 sequences on each sheet. Participants were allowed 10 min in which to study the stimuli (if they asked as to the purpose of the experiment, they were informed that there was to be a memory test). In the test phase (taken by both the experimental and control groups), participants were given response sheets on which the 80 symbol-of-syllable sequences appeared. In all other respects, participants' instructions were identical to those given in Experiment 3.

Results

The percentage correct of classification scores are shown in Table 4.

The data were entered into a two-way ANOVA with variables test set (symbols and syllables) and learning set (none and symbols). There was an effect of test set, $F(1, 22) = 4.6$, $p < .05$, $MSE = 39.50$, an effect of learning set, $F(1, 22) = 34.78$, $p = .0001$, $MSE = 105.00$, but no interaction between the two, $F(1, 22) = 1.67$, $p > .2$, $MSE = 39.50$. Planned comparisons revealed that classification of the syllable test set after exposure to the symbol learning set was significantly better than classification of this set in the control condition, $F(1, 22) = 16.73$, $p = .0005$, $MSE = 105.00$, and that classification of the symbol test set was significantly better than classification of the syllable test set after exposure to the symbol learning set, $F(1, 11) = 10.51$, $p < .01$, $MSE = 39.50$. To determine whether this effect was due to just a few items, we calculated, for each test string, the number of participants giving a correct response in the transfer condition and the corresponding control condition; 22 of the 80 test strings showed significant transfer according to the criterion described in relation to Experiment 3. Given that we would expect by chance alone only four such items, the results again suggest that the transfer effect was not due to only a few items.

Discussion

Experiment 4 has replicated and extended the findings of Experiment 3, and although the underlying grammar used in the two studies was the same, there were a number of differences between the studies (new ungrammatical stimuli and preexposure to symbol sequences instead of syllable sequences) that make the replication significant. The inclusion of a no-transfer condition (and its corresponding control) allows a comparison of no-transfer and transfer classification performance. In the no-transfer and transfer conditions, we can assume that whatever has been extracted on the basis of the learning phase is the same. This knowledge base allowed novel sequences of symbols to be classified at 71% accuracy and sequences of syllables to be classified at a slightly (but significantly) reduced rate of 65%. If we take the figure of 71% as the ceiling against which to compare the magnitude of the transfer effect (given that 71% is a reasonable estimate of the maximum classification performance that we could in principle expect), we find that the absolute level of transfer (15% relative to the control) corresponds to 76% of the ceiling level and that the cost associated with changing the domain of the test stimuli is, correspondingly, 24%.

In the discussion that follows, we consider further the explanations for the transfer effects observed in this and the previous experiments.

Table 4
Percentage Correct of Classification Scores as a Function of Learning Set and Test Set

	Test set			
	Symbols		Syllables	
Learning set	%	*SD*	%	*SD*
Control	51	2.0	49	2.1
Symbols	71	2.5	65	3.0

General Discussion

The results of Experiments 1 to 4 suggest that knowledge acquired during the acquisition phase can be applied to novel domains in the test phase. However, there are a number of possible explanations for this effect. For instance, although the ungrammatical sequences were matched to the grammatical sequences for overall letter–symbol frequencies and length (with an exact match in Experiment 4 and an approximate match in Experiment 3), a proportion of the ungrammatical sequences (between 25% and 35%) started with an element that could not, according to the grammar, occur in initial position (in Experiment 4 it was not possible to eliminate such instances without repeating an ungrammatical sequence or jumbling the elements of the grammatical sequence to create another grammatical sequence). Whereas the grammar and associated vocabulary permitted items starting with either *hes* or *vot*, a proportion of the ungrammatical sequences started with either *jix* or *pel*. In principle, it would be enough to identify these low-frequency starting elements and classify any sequence beginning with such an element as ungrammatical. To rule out such a possibility (which would make the transfer effects we have observed somewhat less interesting), we computed a further post hoc analysis of the data in which we omitted from the analysis the data from sequences with nongrammatical starting elements. Because this resulted in unequal numbers of grammatical and ungrammatical items, we computed d' values[1] for each participant; a measure of the degree to which each participant could successfully discriminate between the grammatical and ungrammatical sequences. A one-way ANOVA comparing d' for the control and transfer groups in Experiments 3 (−0.15 and 0.42, respectively) and 4 (−0.09 and 0.42, respectively) showed that participants' discriminability was significantly improved in the transfer conditions compared with the control conditions—Experiment 3: $F(1, 22) = 10.2, p < .003$, 1-tailed, $MSE = 0.19$, and Experiment 4: $F(1, 21) = 4.65, p < .025$, 1-tailed, $MSE = 0.32$; 1 participant in the transfer group was eliminated because of a missing d'. Of course, these analyses do not rule out the possibility that identification of low-frequency starting elements places an important constraint on any mapping that may be induced across the stimulus set as a whole; low-frequency starting elements should not, for instance, be mapped onto the starting elements identified on the basis of the learning phase. Whether asymmetries in the frequency of occurrence of particular elements play a role in allowing transfer to take place is an open issue, and future research will investigate this issue further. For now, it is important to note that the transfer effects observed here cannot be explained solely in terms of a strategy to reject as ungrammatical sequences beginning with low-frequency elements.

Brooks and Vokey (1991) suggested that in the Reber (1989) and Mathews et al. (1989) studies, transfer between letter strings from different vocabularies could have taken place because of some "abstract analogy" between the letter strings in the test phase and some representation of the individual exemplars in the learning phase. They defined this analogy as "an abstract correspondence of within-item relations among letters" (p. 317). For example, the test sequence *VXVR* can be seen as similar to a learning sequence such as *DFDA* and classified as grammatical without the participants having abstracted a grammar over the learning set. This abstract analogy procedure could be used for each test item separately, without necessarily abstracting over the course of the test phase a mapping between the letter sets (cf. Whittlesea & Dorken, 1993).

To establish whether our data included above-chance performance on sequences without repeats, we computed further post hoc analyses of just the no-repeats data. In Experiment 3, 34 of the sequences (19 grammatical and 15 ungrammatical) contained no repeated elements (whether adjacent or nonadjacent). We computed d' values for each participant, and these were entered into a one-way ANOVA (control vs. transfer) that confirmed that discriminability between grammatical and ungrammatical sequences was significantly greater in the transfer group, 0.27, relative to the control group, −0.17, $F(1, 21) = 4.06$, $p < .03$, 1-tailed, $MSE = 0.28$. Thus, transfer occurred in Experiment 3, even for the subset of sequences containing no repeats. In Experiment 4, there were 38 sequences that contained no repeats (the difference being due to the different ungrammatical stimuli, which were more closely matched in that study). Half of these were grammatical, and half ungrammatical. The equivalent d' analysis (eliminating 2 control participants with indeterminate d's) again confirmed that discriminability between grammatical and ungrammatical sequences was significantly greater in the transfer group, 0.78, relative to the control group, −0.05, $F(1, 20) = 13.7, p < .001$, 1-tailed, $MSE = 0.27$.[2] Of course, these analyses do not rule out the possibility that the mapping is nonetheless induced by exposure to those sequences that do contain repeated items and are then applied to those that do not, but this would still entail that classification is not made simply on an item-by-item basis with no appeal to any underlying knowledge. A test of the hypothesis that repetition structure is necessary to initially induce a mapping would be to ensure that no items, whether in the learning or test phases of the experiment, contained any repeated items. However, this is beyond the scope of the present article, in which the purpose is to demonstrate that transfer can take place, however caused.

Overall, our effects of transfer were relatively small, although we could not expect transfer to be perfect. First, even with no change in domain, learning does not generalize perfectly across the learning and test sets (and nor can it, given that only a subset of the exemplars generated by the grammar are ever presented during the learning phase), and thus an upper limit is set on performance when the learning and test sets are in different domains. Second, transfer cannot be perfect because of the noise introduced by the ungrammatical exemplars present in the test set. Given that even without a change in domain, the system is not perfectly sensitive to which sequences are grammatical and which are not, it follows that a mapping will be induced on the basis of both grammatical and ungrammatical items, even though only the grammatical items "define" the appropriate mapping(s). Finally, on the assumption that the computation of the relevant mappings takes time to develop during exposure to successive exemplars, it would take some time before correct classification of novel stimuli in a novel domain could take place (see the *Results* section of Experiment 1 for further discussion and the computational simulation described below for a computational instantiation of this assumption).

Our results pose a challenge for the various accounts of artificial grammar learning (e.g., Brooks, 1978; Cleeremans & McClelland, 1991; Dienes, 1992; Druhan & Mathews, 1989; Roussel & Mathews, 1994; Perruchet & Pacteau, 1990; Servan-Schreiber & Anderson, 1990; Vokey & Brooks, 1992; see Berry & Dienes, 1993, for a review). Brooks pointed out that participants could learn an artificial grammar by storing representations of each training string and then by classifying test strings according to their similarity to the stored training strings. Brooks and Vokey (1991) showed further how such stored strings could allow transfer between different domains by the process of abstract analogy we described previously. The classifier system described by Druhan and Mathews and by

[1] $$d' = \frac{\sqrt{3} \times \ln\left(\frac{\text{hits} \times \text{correct rejections}}{\text{misses} \times \text{false alarms}}\right)}{\pi}$$

[2] Only 8 of the 50 test sequences from Experiment 1 contained no repeats, and there were too many missing d' values to compute an analysis. However, because there were equal numbers of grammatical and ungrammatical sequences, we computed a two-way ANOVA to test the transfer effect across the two test types (letters and tones) for just the eight items without repeated elements. The main effect of transfer (i.e., control vs. transfer) was significant, $F(1, 22) = 4.38, p < .05$, $MSE = 0.03$, and there was no main effect of, or interaction with, test type (both Fs < 1).

Roussel and Mathews produced transfer in a very similar way. In their system, the strength of features of exemplars was tuned according to its ability to accurately predict grammaticality. Crucially, the features could be more than specific letter sequences, they could also be an abstract pattern of adjacent repeats (e.g., *MTTTV*, could be encoded as __rr_, where *r'* stands for a repeat of the immediately preceding letter). That is, the classifier system allowed transfer only in so far as the different domains had similar patterns of adjacent repeats (sometimes referred to as *runs*) in each string. In this respect it is similar to the Brooks and Vokey proposal, except that for Brooks and Vokey, nonadjacent repeats would also be an example of the kinds of within-item relations among letters that could give rise to transfer (and if there were no repeats anywhere, all the letters would of course be different, and there could be no within-item relations across sequences drawn from the two domains that could correspond, unless they did so by virtue of some form of abstract knowledge equivalent to an underlying grammar). Our data indicate that this cannot be the only mechanism: In the new domain, participants could correctly classify strings in which there were no repeats, either adjacent or separated. Participants must have induced at least a partial mapping between the domains (and see our earlier discussion of this point).

Perruchet and Pacteau (1990) argued that participants may use stored fragments of exemplars, especially bigrams, as a basis for grammaticality judgments. However, Gomez and Schvaneveldt (1994) and Manza and Reber (1994) showed that knowledge of bigrams could not underlie transfer between domains because participants trained only on bigrams could not classify strings in a new domain at above-chance levels. Of course, participants may build up representations of commonly occurring higher order fragments (Servan-Schreiber & Anderson, 1990), but such knowledge remains domain specific, and the problem is how to induce a mapping between fragments in one domain and stimuli in another. Similarly, the connectionist models used by Dienes (1992) and the simple recurrent network (SRN) used by Cleeremans and McClelland (1991) create representations tied to the particular domain in which the network was trained, and an additional mechanism is needed to effect a mapping.

We (Dienes, Altmann, & Gao, 1994, 1995) have recently shown how a variation of an SRN, when given exactly the same training strings used in the experiments described above, can classify test strings at levels of accuracy comparable with human participants. The architecture we used is shown in Figure 4.

The input layer is divided into two parts: D1 for coding the information in the first domain, and D2 for coding the information in the second domain. The first layer of hidden units then recodes both domains, and this recoding is used as an input to a standard SRN, with a separate D1 and D2 output layer. All weights are initialized as small random values. The first element of a sequence in, for instance, the symbol domain, would be coded by the D1 input units. This is recoded by the hidden layers to predict the second symbol of the sequence. Weights are adjusted by backpropagation. Then the second symbol is applied to the D1 input units and so on. In a subsequent test phase, the network more successfully predicts successive symbols of grammatical rather than nongrammatical sequences when they are applied to D1, and this fact can be used to produce equivalent same-domain (i.e., no-transfer) classification performance as people for equivalent training (i.e., four epochs). When testing the network in a new domain (e.g., syllable sequences), the recurrent weights and the weights between the hidden layers (the core weights in Figure 4) are frozen, or their learning rates reduced (as might happen following the application of an optimization rule that reduces the learning rate as a function of the sign and magnitude of the backpropagated error), and only the D2 input and output mapping weights (see Figure 4) are changed. The D2 mapping weights start at arbitrary random values. The first syllable is applied to the D2 input units, and the network attempts to predict the second syllable. Backpropagation changes the mapping weights, and the network then attempts to predict the third syllable given the second and so on. By the time the network has reached the end of the sequence, the mapping weights have changed, and the network can iterate around the sequence a number of times before moving on to the next test sequence. As in the no-transfer case, the network will classify a string as grammatical if on the last iteration around the sequence it can predict successive syllables in the sequence well.

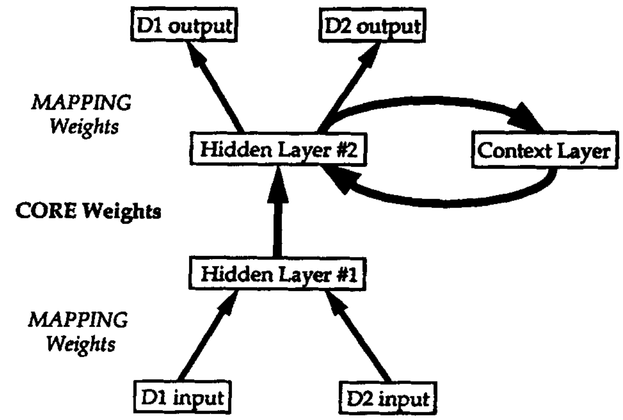

Figure 4. The architecture from Dienes, Altmann, and Gao, 1994. D1 = first part of the layer used for coding the information in the first domain; D2 = second part of the layer used for coding the information in the second domain.

Because the weights between the hidden layers implicitly encode the structure of the new domain (given that this same structure has been encoded on the basis of exposure to sequences from the original domain), the network just needs to learn the mapping to and from the abstract encodings formed by the SRN. Despite the noise introduced by adjusting the mapping weights when presented with nongrammatical strings, the network does indeed learn the mappings. There is an advantage of training the network on the same grammar as opposed to a different grammar to that of the test stimuli. When the network is trained to produce equivalent same-domain performance as people, it can produce equivalent cross-domain transfer. If the same-domain performance is taken to define the maximum amount of cross-domain transfer that could in principle be shown (cf. discussion of the results of Experiments 1 and 4), then the model, like people, can perform at about 70% of the maximum possible. Analysis of the weights to the first hidden layer shows how the correct mapping between the domains is gradually and partially induced over the course of the test sequences (and this is discussed further in Dienes et al., 1995).

Our earlier assumptions about the limitations on learning are supported by the computational simulations just described. The model allowed us to separate out the contribution to performance that was due to how much was learned on the basis of the training exemplars and the contribution that was due to how good a mapping is derived during the test phase. We do not believe that the SRN is necessarily the best model for the data we have, but it does provide useful insights into what information can be extracted and how such extraction may be influenced by properties of the stimulus set (see also Cleeremans, Servan-Schreiber, & McClelland, 1989; Dienes, 1992).

Despite the support provided by the computational simulations, there is no guarantee that the way in which the network solves the mapping problem is the same as the way in which people do. Whittlesea and Dorken (1993) argued that the acquired knowledge that underlies the ability to transfer across stimuli sets is not abstract knowledge that is induced across the item set (similar in respects to Brooks and Vokey, 1991). They described an experiment (Experiment 4) that "had all the characteristics of a standard test of the abstractness of implicit grammar learning, except that it had no grammar" (Whittlesea & Dorken, 1993, p. 238). That is, there was "no description of the [training] set in terms of typical members or a set of rules smaller than the set of [individual] instances" (Whittlesea & Dorken, 1993, p. 238). They found that, nonetheless, participants could discriminate between test sequences that had been mapped onto a new letter set and novel sequences that apparently obeyed the same "grammarless" property as the test set. They argued that the "grammarlessness" of the stimuli prevented the acquisition of any knowledge that could have been induced across the training set and that participants' performance must have been due to knowledge about (the deep-structural properties of) individual items. However, inspection of Whittlesea and Dorken's materials reveals that, in principle, a single rule could be induced across the training set that would allow discrimination between their "legal" and "illegal" test items—their stimuli were of the form "1234–2413" or "1231–2443," where these patterns describe the repetition pattern of actual stimuli such as *PTZC–TCPZ* and *PGTP–GVVT*. A mean accuracy of 0.65 could be achieved in the test phase by classifying as ungrammatical any sequence in which a binary transition in the first half of the stimulus recurred in the second half of the stimulus (e.g., *VTCK–TCVK*). Only 3 of the 16 training items contained such recurrences, whereas 8 of the 16 "new" test items did. Although this rule would not perfectly describe the training set, its application would in fact lead to greater accuracy than that reported for the participants in the experiment (0.57). We believe, therefore, that participants could in principle have induced abstract knowledge across the relevant stimuli sets in that experiment (although whether they would have abstracted the rule we have identified or some different rule(s) is unclear), and, consequently, Whittlesea and Dorken's results do not contradict our own.

The four studies described above demonstrate that a grammar acquired in one domain can be used to categorize structures in another domain, albeit imperfectly. That is, regularities extracted from within one domain can be used to impose order on the regularities that might occur in another domain. Although the experiments we have so far described have confounded domain with modality (see our earlier definition of what constitutes a change in domain), we believe that changing modality and domain constitutes the strongest possible evidence for transfer effects. One immediate interpretation of our data (we offer another interpretation below) is that however the regularities are internally represented, they are represented in a form that is domain or modality independent. To the extent that we might define such internal representations as constituting a grammar, then it follows (according to this interpretation) that the grammar itself is available to processes operating in domains as diverse as the auditory and visual domains. Although such a hypothesis possibly constrains claims about what might or might not constitute Fodorian modules (cf. Fodor, 1983) or where, in a Fodorian architecture, learning of the kind described here takes place, more important is the demonstration that the human processing system does permit, one way or another, the cross-domain transfer of learned grammatical structure. Whether this transfer is brought about by the construction of a domain-independent grammar, as opposed to a direct mapping between tokens in one domain and some internal representation of tokens in another, is uncertain.

One issue that remains outstanding is whether the learning that took place in Experiments 1 to 4 was implicit and whether the ensuing transfer that was also observed in those experiments was itself implicit. We can present no data that determine, one way or the other, the implicitness of the knowledge that was abstracted on the basis of the test set of exemplars. Nonetheless, our results warrant further investigation in this regard, and we are currently exploring the extent to which the transfer effects we have observed are indeed implicit.

A related issue concerns the extent to which implicit memory may be implicated in transfer effects. The finding that knowledge of an artificial grammar can transfer across modalities apparently contrasts with the typical finding in the implicit memory literature (see Schacter, 1987, for review). For example, if participants study pictures rather than words, fragment completion of the corresponding words is markedly reduced (Weldon & Roediger, 1987). Furthermore, priming (i.e., facilitation of fragment completion on the basis of the prior study period) is impaired even if words are used at study and test, but there is a shift in modality (Bassili, Smith, & MacLeod, 1989). The cost of changing modalities in Experiments 1 (38%) and 4 (24%) is similar to the decrement found in the implicit memory literature (as overviewed by Dienes & Fahey, 1994), although Experiments 1 and 2 involved not just a change in modality but also a random mapping between the modalities. Further research needs to explore transfer in artificial grammar learning with a transparent mapping to establish whether the decrement found there is of the same magnitude as that found in the implicit memory literature (and so establish whether the decrement we observe here is primarily due to the change in modality, the random mapping, or both). One relevant difference between artificial grammar learning and implicit memory tasks may be that the knowledge in artificial grammar learning is entirely about the establishment of new associations; in implicit memory, the emphasis is on the reconstruction of stimuli that are previously well known (this is true also of the phenomena of implicit memory for new associations; Schacter & Graf, 1989). In summary, although we did not set out to consider issues in implicit memory, our transfer data suggest a possible dissociation between transfer effects found in the study of implicit memory and transfer effects found in the study of (implicit) learning (see Berry and Dienes, 1991, for further discussion of the relation between artificial grammar learning and implicit memory).

Notwithstanding the largely unresolved issue concerning the implicitness of the knowledge that was acquired, and applied, in our studies, it is pertinent to consider further the processes that may underlie implicit learning and the relationship that these processes may have to the transfer effects we (and others) have observed. The result of implicit grammar learning must be some internal representation onto which the external input can be projected or mapped. In effect, this mapping process is a form of parsing. This view of the process is agnostic as to whether the internal representation is some analogue to a (phrase-structure or finite-state) grammar or whether the representation is in terms of distributional characteristics of word types (cf. syntactic form classes) or even word tokens. What matters is that some mapping is performed between the external input and this internal representation. For example, the external input needs to be mapped onto the set of fragments or chunks extracted from the training stimuli (Servan-Schreiber & Anderson, 1990), or onto the appropriate hidden unit representations used for recoding the training stimuli (Cleeremans & McClelland, 1991), or onto whatever underlying abstract grammar has been induced through exposure to the training stimuli. In the case of a classification task that uses the same vocabulary as that on which basis the grammar was acquired, the processor attempts to parse the new input, and if successful the input is judged grammatical, and if not, the input is judged ungrammatical. The challenge is to explain the processing that accompanies cases in which the classification task uses a different vocabulary from that on which basis the grammar was acquired. These cases include not simply our own but also other demonstrations of transfer between different letter sets (e.g., Brooks & Vokey, 1991; Mathews et al., 1989; Reber, 1969; Whittlesea & Dorken, 1993). In these cases, the processor must establish a mapping between the new vocabulary and the old vocabulary (or underlying form classes) as it attempts to parse the new input sequences. Further research will investigate the stability of this mapping, given that it need not be absolutely stable but could instead be probabilistic.

While it might be tempting, at least at first, to propose the existence of a domain-independent representation of grammatical knowledge, an alternative proposal would consist of a domain-dependent representation with the transfer effects that we observed being due to domain-independent processes that can operate across domains or modalities (cf. the analogical process advocated by Brooks and Vokey, 1991, discussed above). In either case, some form of mapping is required, and in either case, knowledge about regularities in one domain is both available and, we suggest, automatically applied to the task of categorizing novel stimuli in another domain.

Although the artificial grammar learning task described here has involved the presentation of only relatively simple stimuli, sensitivity to quite complex linguistic structures can implicitly be acquired, even when these structures are embedded in natural speech stimuli. For instance, Zwitserlood (1990) reported a study in which Dutch adults who had no knowledge of Mandarin Chinese were exposed to Mandarin speech for just 12 min. The speech was presented by a cartoon film narrated by a native Mandarin speaker. Participants were not informed of their task or what language they were listening to. Subsequently, participants were presented with a sequence of spoken stimuli consisting of real Mandarin words interspersed with pseudowords that violated a variety of phonotactic constraints in Mandarin (e.g., segments appeared in the wrong order within the word–syllable). Zwitserlood found that even with such brief exposure, participants could discriminate between the real words and the pseudowords at levels that were well above chance. This study demonstrated that phonotactic constraints (which constitute one of a range of linguistic structures) can be acquired very quickly on the basis of very brief exposure. Thus, exposure to complex stimuli (i.e., naturally occurring speech) can induce implicit learning and, subsequently, sensitivity to constraints on the sequentiality of natural language structures. It is in the context of modeling the acquisition of constraints on sequentiality that there has been growing interest in artificial neural networks that, although modeling processes identical to those implicated in implicit learning as studied in the laboratory (cf. models described by Berry & Dienes, 1993; Cleeremans et al., 1989; Dienes, 1992; Dienes et al., 1994, 1995), have been shown to acquire operating characteristics considered desirable in any model of the acquisition of language (cf. Elman, 1990a, 1990b, 1993, and the emergence of syntactic categories and sensitivity to grammatical structure; Cleeremans et al., 1989, and Sopena, 1991, and the emergence of sensitivity to long-distance dependencies and other linguistic phenomena; and Norris, 1990, and the ability to normalize in the temporal domain).

Despite some potential (but highly controversial) linkage between implicit learning and natural language (see Morgan et al., 1987; Winter & Reber, in press), it is unclear whether the transfer effects we have described could themselves be at all relevant to natural language.[3] Before any such speculation is possible, it must first be determined whether the ability to transfer across domains is only limited by the capacity to

[3] There do exist cases in natural language in which the same underlying grammar generates sequences in two different domains. In Arabic, for instance, the sound-to-spelling rules are totally regular, with a one-to-one correspondence between spelling and pronunciation. Spanish is an example of another regular language, and even the English alphabet is largely phonetic. It would be feasible to determine whether a degree of transfer is possible across these domains by using a technique similar to that used by Zwitserlood (1990), with nonspeakers of Arabic listening to Arabic words, but with orthographic stimuli presented at test. However, even if transfer did occur, this would not mean that it is implicated in the normal acquisition of orthography or in language acquisition more generally.

abstract structure implicitly in either domain or whether there might be other limitations on transfer.

References

Bassili, J. N., Smith, M. C., & MacLeod, C. M. (1989). Auditory and visual word stem completion: Separating data-driven and conceptually-driven processes. *Quarterly Journal of Experimental Psychology, 41,* 439–453.

Berry, D., & Dienes, Z. (1991). The relationship between implicit memory and implicit learning. *British Journal of Psychology, 82,* 359–373.

Berry, D., & Dienes, Z. (1993). *Implicit learning.* Hillsdale, NJ: Erlbaum.

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

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.

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

Cleeremans, A., Servan-Schreiber, D., & McClelland, J. L. (1989). Finite state automata and simple recurrent networks. *Neural Computation, 1,* 372–381.

Dienes, Z. (1992). Connectionist and memory-array models of artificial grammar learning. *Cognitive Science, 16,* 41–79.

Dienes, Z., Altmann, G. T. M., & Gao, S. (1994, April). The transfer of implicit knowledge across domains. Poster presented at the fourth meeting of the Joint Council Initiative on Human Computer Interaction and Cognitive Science. Abstract to appear in *Language and Cognitive Processes.*

Dienes, Z., Altmann, G. T. M., & Gao, S. (1995). Mapping across domains without feedback: A neural network model of implicit knowledge. In L. S. Smith & P. J. B. Hancock (Eds.), *Neural computation and psychology: Proceedings at the 3rd Neural Computation and Psychology Workshop (NCPWS), Stirling, Scotland, 31 August–2 September, 1994.* Springer-Verlag: London.

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.

Dienes, Z., & Fahey, R. (1994). *The role of implicit memory in controlling a complex system.* Unpublished manuscript.

Druhan, B., & Mathews, R. (1989). THIYOS: A classifier system model of implicit knowledge of artificial grammars. *Proceedings of the 11th annual conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

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

Elman, J. (1990a). Representation and structure in connectionist models. In G. T. M. Altmann (Ed.), *Cognitive models of speech processing: Psycholinguistic and computational perspectives* (pp. 345–382). Cambridge, MA: MIT Press/Bradford Books.

Elman, J. (1990b). Finding structure in time. *Cognitive Science, 14,* 179–211.

Elman, J. (1993). Learning and development in neural networks: The importance of starting small. *Cognition, 48*(1), 71–99.

Fodor, J. A. (1983). *Modularity of mind.* Cambridge, MA: MIT Press.

Gomez, R. L., & Schvaneveldt, R. W. (1994). What is learned from artificial grammars?: Transfer tests of simple association. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 396–410.

Manza, L., & Reber, A. S. (1991). *Implicit learning: Transfer across form and sensory modality.* Poster session presented at the third annual meeting of the American Psychological Society, Washington, DC.

Manza, L., & Reber, A. S. (1994). *The representation of tacit knowledge: II. Transfer between different letter sets.* Manuscript in preparation.

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

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

Morgan, J. L., Meier, R. P., & Newport, E. L. (1987). Structural packaging in the input to language learning: Contributions of prosodic and morphological marking of phrases to the acquisition of language. *Cognitive Psychology, 19,* 498–550.

Norris, D. (1990). A dynamic-net model of human speech recognition. In G. T. M. Altmann (Ed.), *Cognitive models of speech processing: Psycholinguistic and computational perspectives* (pp. 87–104). Cambridge, MA: MIT Press/Bradford Books.

Perruchet, P. (1994). Learning from complex rule-governed environments: On the proper function of unconscious and conscious processes. In C. Umiltà & M. Moscovitch (Eds.), *Attention and Performance XV: Conscious and nonconscious information processing.* Cambridge, MA: MIT Press.

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 Behaviour, 77,* 317–327.

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

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

Roussel, L., & Mathews, R. (1994). *THIYOS: A synthesis of rule-based and exemplar-based models of implicit learning.* Unpublished manuscript.

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

Schacter, D., & Graf, P. (1989). Modality specificity of implicit memory for new associations. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 15,* 3–12.

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

Sopena, J. M. (1991). *ERSP: A distributed connectionist parser that uses embedded sequences to represent structure.* (Tech. Rep. No. UB-PB-1-91). Departament de Psicologia Bàsica, University of Barcelona.

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

Weldon, M. S., & Roediger, H. L. (1987). Altering retrieval demands reverses the picture superiority effect. *Memory & Cognition, 15,* 269–280.

Whittlesea, B. W. A., & Dorken, M. D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General, 122,* 227–248.

Winter, B., & Reber, A. S. (in press). Implicit learning and natural language acquisition. In N. Ellis (Ed.), *Implicit and explicit learning of languages.* London: Academic Press.

Zwitserlood, P. (1990). *Speech segmentation during the first minutes of second-language acquisition.* Paper presented to the 31st annual meeting of the Psychonomic Society, New Orleans, LA.

Appendix A

Training Items Used in Experiments 3 and 4

In Experiment 3, these items were spoken with intonation consistent with phrase boundaries as indicated with |. For Experiment 4, each of these sequences was translated into a sequence of symbols according to the mapping given in Table 3 (the | signs were not included). The order shown below is for the first block of the 30 training items. The blocks were all presented in different random orders.

hes jix | pel jix dup
hes | pel kav | pel
hes | pel dup
hes jix | rud | pel
hes jix | pel jix dup | pel
vot jix | pel jix kav
hes | pel jix dup
vot | pel jix kav
hes | pel jix kav | pel
hes | sog
vot | sog | pel
hes | pel dup | pel jix
vot jix | pel jix kav | pel
hes jix | sog
vot | pel jix dup | pel
vot jix | pel dup | pel jix |
vot | pel kav
hes | rud
vot jix | rud
hes | pel dup | pel
hes jix | pel kav | pel jix
vot | pel kav | pel jix
hes jix | pel jix dup | pel jix
vot jix | sog | pel jix
vot | rud | pel jix
vot | pel kav | pel
vot jix | sog | pel
hes | pel jix kav | pel jix
hes jix | rud | pel jix
vot | pel jix dup | pel jix

Appendix B

Test Items Used in Experiment 3

Ungrammatical sequences are marked with an asterisk. Items marked with a dagger showed significant transfer according to the criterion discussed in the main text.

vot jix pel dup
hes jix hes pel sog*
hes jix sog pel jix
pel jix dup hes pel jix*
vot kav jix*
vot jix pel kav pel jix
kav pel jix kav†*
vot pel rud pel jix pel*
hes jix pel jix kav
kav jix rud*
vot jix vot jix dup pel†*
vot jix pel kav pel
hes pel kav†
hes jix hes kav pel jix kav†*
vot jix pel jix dup pel
hes kav pel*
vot sog pel jix
jix vot jix pel kav*
hes pel jix dup*
hes pel kav pel kav†*
hes jix pel jix kav pel jix
vot hes vot dup*
hes pel jix kav
hes jix rud
vot pel dup pel jix†
hes jix vot jix kav†*
jix rud dup pel*
hes pel jix pel jix†*
hes pel jix dup pel†
vot sog
vot jix dup kav dup†*
hes pel jix dup pel jix
hes jix pel dup
vot jix hes pel*
hes jix pel kav pel
vot jix rud pel†
vot pel dup pel
vot dup*
vot hes vot hes†*
vot rud sog
dup jix sog*
vot jix pel kav
hes jix vot kav*
vot jix sog†
hes jix pel jix kav pel
vot pel dup†
hes pel kav pel jix
hes jix hes jix kav pel†*
hes sog pel jix
vot jix hes jix dup*
dup pel jix dup†*
hes kav*
hes jix pel dup pel
hes sog pel†
vot jix vot dup pel jix dup*
vot jix rud pel jix
hes vot hes pel*
vot jix pel jix dup
vot pel jix pel jix*
vot dup pel*
vot pel jix kav pel jix
jix vot kav pel*
hes dup jix pel*
vot pel dup pel dup†*
hes jix pel kav
hes pel sog pel jix pel†*
hes jix sog pel
vot pel jix kav pel
hes vot hes vot†*
hes jix kav dup kav*
hes dup jix*
hes rud
vot jix pel dup pel
hes rud pel jix
jix vot jix pel dup*
vot jix pel jix dup pel jix
pel jix kav vot pel jix†*
vot jix vot pel rud*
hes jix pel dup pel jix
vot pel jix dup

(*Appendix C follows on next page*)

Appendix C

Test Items Used in Experiment 4

Ungrammatical sequences are marked with an asterisk. These same sequences were used, in the reverse order, for testing on symbols, in which case they were translated into symbol sequences according to the mapping shown in Figure 3. Items marked with a dagger showed significant transfer according to the criterion discussed in the main text.

jix hes pel dup pel†*
pel vot jix dup pel†*
vot jix sog†
hes kav jix pel*
hes sog pel jix
vot pel dup pel jix
hes jix dup pel jix pel*
vot pel kav jix pel*
hes jix pel jix kav
vot dup*
vot jix pel dup
vot kav pel jix pel*
hes jix kav pel jix pel*
hes rud
vot jix pel kav pel
hes pel kav jix pel*
jix hes rud*
vot sog†
hes pel jix sog*
jix vot dup jix pel†*
vot pel jix kav pel
pel hes kav†*
hes rud pel jix
vot pel jix dup†
pel hes jix kav jix pel jix†*
hes jix pel kav
hes pel kav

vot pel sog jix*
hes sog pel†
vot jix rud pel jix
hes dup pel jix pel*
vot pel rud jix*
hes jix pel dup pel jix
hes kav pel jix pel*
hes jix pel sog jix*
hes jix pel kav pel
jix hes kav pel†*
hes pel kav pel jix†
jix vot pel rud jix†*
vot jix pel jix dup pel
pel hes sog*
pel vot jix kav pel jix†*
vot jix pel jix dup pel jix
hes pel sog jix*
vot jix pel jix dup
vot pel jix kav pel jix
vot dup pel jix pel†*
vot pel dup pel†
hes jix pel jix kav pel jix
vot sog pel jix
vot pel rud*
hes jix pel jix kav pel
vot jix rud pel
jix vot kav pel†*

hes jix rud†
hes jix sog pel jix
hes dup jix pel jix pel†*
jix hes dup pel*
hes jix pel dup
vot jix dup jix pel jix pel*
vot sog jix*
vot jix pel kav pel jix†
pel vot dup jix†*
vot jix dup pel*
hes jix pel dup pel
hes kav*
pel vot sog pel*
hes pel jix dup pel
vot dup pel*
hes jix sog pel†
vot rud pel
hes jix rud jix pel*
hes pel jix dup pel jix
vot jix pel dup pel
hes pel rud jix*
vot jix kav pel jix pel*
hes pel jix kav
vot dup jix pel jix pel*
vot jix pel kav
vot pel dup†

Received February 28, 1994
Revision received July 11, 1994
Accepted August 31, 1994 ■