The basis of transfer in artificial grammar learning

REBECCA L. GOMEZ and LOUANN GERKEN
University of Arizona, Tucson, Arizona

and

ROGER W. SCHVANEVELDT
New Mexico State University, Las Cruces, New Mexico

In two experiments, we examined the extent to which knowledge of sequential dependencies and/or patterns of repeating elements is used during transfer in artificial grammar learning. According to one view of transfer, learners abstract the grammar's sequential dependencies and then learn a mapping to new vocabulary at test (Dienes, Altmann, & Gao, 1999). Elements that are repeated have no special status on this view, and so a logical prediction is that learners should transfer as well after exposure to a grammar without repetitions as after exposure to a grammar with them. On another view, repetition structure is the very basis of transfer (Brooks & Vokey, 1991; Mathews & Roussel, 1997). Learners were trained on grammars with or without repeating elements to test these competing views. Learners demonstrated considerable knowledge of sequential dependencies in their training vocabulary but did not use such knowledge to transfer to a new vocabulary. Transfer only occurred in the presence of repetition structure, demonstrating this to be the basis of transfer.

Artificial grammar learning (Reber, 1967) has been used extensively to study how learners acquire complex knowledge. In a typical experiment, learners are exposed to a subset of exemplars generated by a finite-state grammar (see Figure 1). Participants are trained in the context of a cover task (e.g., observation, typing, or memorization to criterion), so that any learning is presumably incidental in nature. After acquisition, participants are given a surprise test, in which they are asked to discriminate new valid exemplars from invalid ones. Learning occurs rapidly (see, e.g., Altmann, Dienes, & Goode, 1995), occurs under a number of training conditions (e.g., Gomez & Schvaneveldt, 1994; Mathews et al., 1989; Reber, 1969; Reber & Allen, 1978), and is retained over as much as a 2-year delay (Allen & Reber, 1980). It occurs both early (Gomez & Gerken, 1996, 1999) and late (Meulemans & Van der Linden, 1997b) in development and in a variety of patient populations, including those with psychological and/or neurological impairment (see, e.g., Abrams & Reber, 1988; Knowlton & Squire, 1994, 1996).

One of the most remarkable findings in artificial grammar learning is transfer to new vocabulary. During transfer, the elements (or vocabulary) from which training strings are composed are mapped onto new vocabulary, so that the underlying syntax is identical in training and test but the surface vocabulary is changed. Transfer has been documented extensively in adults (Altmann et al.,

1995; Brooks & Vokey, 1991; Gomez, 1997; Gomez & Schvaneveldt, 1994; Knowlton & Squire, 1996; Mathews et al., 1989; Reber, 1969; Vokey & Higham, 1997; Whittlesea & Dorken, 1993; Whittlesea & Wright, 1997) and, most recently, in learners as young as 7 (Marcus, Vijayan, Bandi Rao, & Vishton, 1999) and 11 months of age (Gomez & Gerken, 1997, 1999). Such findings are crucial for demonstrating that learners are not only acquiring specific sequential dependencies but also are abstracting some aspect of the sequential structure. In Reber's own words, "it seems clear from these results that Ss are learning abstract structures to which a particular set of symbols are assigned. They are not simply learning to string together explicit symbols" (Reber, 1969, p. 119).

Although vocabulary-same performance is fairly well understood, we know little about the mechanisms underlying transfer. For instance, learners tested in their training vocabulary obtain knowledge of the associative strength of legal bigrams and trigrams (Knowlton & Squire, 1994; Meulemans & Van der Linden, 1997a). They also obtain considerable knowledge of first- and second-order sequential dependencies, where first-order dependencies are determined by the immediately preceding element and second-order dependencies are determined by the preceding two elements (Dienes, Broadbent, & Berry, 1991; Gomez, 1997; Gomez & Schvaneveldt, 1994; Perruchet & Pacteau, 1990). We know that knowledge of first-order dependencies does not produce transfer. Learners trained only on legal bigrams detected violations of first-order dependencies in their training vocabulary, but not during transfer (Gomez & Schvaneveldt, 1994). The same learners were unable to detect violations of second-order dependencies. In contrast, learners trained on entire strings detected violations of second-order dependencies and also
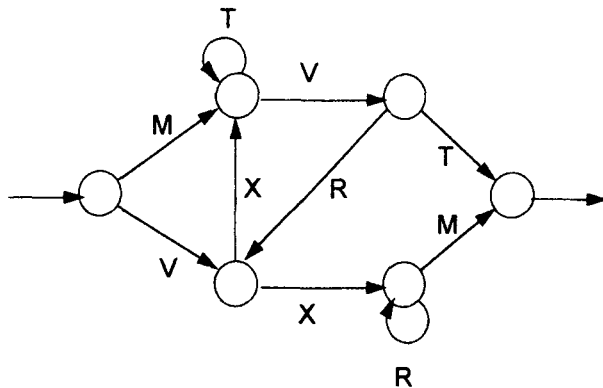
Figure 1. Finite-state grammar used in a number of artificial grammar learning studies (e.g., Altmann, Dienes, & Goode, 1995; Dienes, Broadbent, & Berry, 1991, Reber & Allen, 1978). Grammatical strings are produced by starting at the leftmost state and then traversing the links in the grammar from left to right. Examples of grammatical strings are MTTTVT, VXVT, VXVRXM, and VXRRRR.

transferred to new vocabulary, suggesting that transfer reflects the abstraction of second-order dependencies. However, a salient feature of most artificial grammars is the prevalence of repeating elements (see Figure 1). Cycles are permitted, and elements occur in more than one position in the grammar, resulting in patterns of relations over perceptually identical elements (or repetition structure; e.g., MTTVT, VXVRXRM), raising the issue of whether learners transfer on the basis of knowledge of second-order dependencies, repetition structure, or both.

Two general views have been advanced to explain transfer in artificial grammar learning. Sensitivity to repetition structure is essential to one view (Brooks & Vokey, 1991; Mathews & Roussel, 1997) and has no special status in the other (Altmann et al., 1995; Dienes, Altmann, & Gao, 1999).

According to Mathews and Roussel (1997), learners acquire condition–action rules for encoding letter patterns by position. Given the training string WWSNPZ, learners might induce the rule *strings beginning with a repeating element are grammatical*. Although the notation used here differs somewhat from that used by Mathews and Roussel, such a rule could be encoded as xx_ _ _ _ (to represent any string beginning with a repeat of the first element). Given the string WSWNPZ, learners might induce a rule allowing repetitions that span a single element. Such a rule could be encoded as x_x_ _ _.[1] According to this model, the strength of individual rules increases with the number of matches to test strings, and therefore, the resulting learning reflects aggregated knowledge of the grammar.

According to exemplar theory (Brooks & Vokey, 1991), learners store individual training instances and then compute the similarity between new and stored exemplars at test. The similarity comparisons executed during transfer capitalize on patterns of relations among repeating el-

ements in new and stored exemplars. A match between the training instance BBGXTR and the test string WWSNPZ would result in a judgment of grammaticality, because of the identical repetition structure.

Although repetition structure is central to both Mathews and Roussel (1997) and Brooks and Vokey (1991), it is important to note that the mechanism by which transfer is accomplished differs on these accounts. According to Mathews and Roussel, transfer is based on aggregated knowledge, whereas Brooks and Vokey argue that transfer results from analogy to specific training instances. Nevertheless, a prediction consistent with both views is that transfer should occur after exposure to a grammar with repetition structure but not after exposure to a grammar with no repeating elements. Such transfer could be taken as evidence of the ability to abstract patterns of relations among perceptually identical elements, but not the grammar's sequential dependencies.

According to a very different view (Altmann et al., 1995; Dienes et al., 1999), learners abstract the structure of the grammar from sequential dependencies in the input and then induce a mapping to new vocabulary at test. Dienes et al. (1999) tested such transfer, using the model depicted in Figure 2. The initial layers (V1 and V2) code input in either of two domains (or vocabularies). The first hidden layer (the hidden encoding layer) recodes the initial input in a form independent of specific vocabulary. The recoded input is then fed to a simple recurrent network (SRN; Elman, 1990). The SRN learns sequential dependencies by means of a context layer that holds a copy of the hidden layer's activation of units at the previous time slice. When processing an element in a sequence, activation is passed from the input units to the hidden encoding layer. Input from this layer and the context layer is fed through the hidden units of the SRN to predict the next element in the sequence. Weights are then adjusted, using backpropagation.

During transfer, the core weights are frozen so that only the mapping weights change. The network's task is then to learn a mapping from the SRN's abstract encoding to new vocabulary with feedback obtained from predicting successive elements in sequence. The mapping weights are adjusted (via backpropagation) after each prediction. The network's accuracy in predicting successive elements is used during classification, so that high accuracy leads to "grammatical" judgments and low accuracy leads to "nongrammatical" ones.

Repetition structure has no special status by this view— that is, in principle, sequential dependencies between repeated and nonrepeated elements should not be learned any differently. Therefore, it is reasonable to predict that transfer should occur as readily after exposure to grammars without repetitions as it does after exposure to grammars with them.

Consistent with this view, Altmann et al. (1995) reported post hoc tests reflecting significant discrimination for items with no repetitions. However, a number of
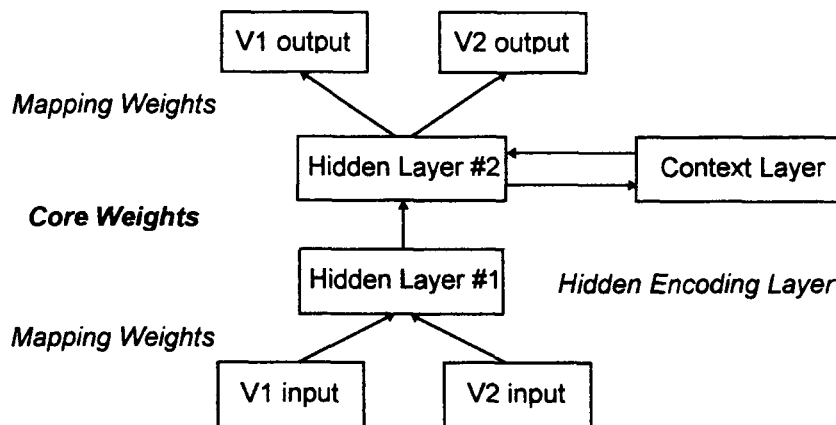
Figure 2. Learners abstract the structure of the grammar via sequential processing of input, then induce a mapping to new vocabulary (see Altmann, Dienes, & Goode, 1995; Dienes, Altmann, & Gao, 1999; also Elman, 1990).

the ungrammatical strings used in these comparisons began with illegal elements. Violations in string-initial position are particularly salient for learners (Reber & Allen, 1978; Reber & Lewis, 1977), raising the possibility that the Altmann et al. analysis reflected knowledge of the frequency distributions of initial elements, rather than transfer based on knowledge of grammatical structure (see Tunney & Altmann, 1999). In contrast, repeating elements appear to be particularly salient for learners. For instance, Tunney and Altmann reported that learners cannot discriminate legal and illegal strings during transfer when grammatical repetition structure is preserved in illegal items. Furthermore, Shanks, Johnstone, & Staggs (1997) investigated detection of various violation types during transfer and found that learners showed reliable detection of illegal repetitions. Finally, Vokey and Higham (1997) found that the tendency to judge transfer strings to be similar to training strings increased with an increase in the number of repetitions.

In summary, several studies attest to the salience of repetition structure during transfer, but no study has adequately tested the possibility that learners can abstract the structure of the grammar in the absence of repetitions. The crux of the matter is this: If learners can transfer without repetitions, this tells us that they are abstracting fairly complex sequential information (involving, at a minimum, second-order dependencies between elements in sequence, as in Gomez & Schvaneveldt, 1994). Such transfer bears similarity to a task required of natural language users—namely, abstraction of the sequential dependencies between syntactic constituents. In contrast, transfer based on repetition structure indicates abstraction but bears little resemblance to the abstraction abilities implicated in natural language use. Rather, such abstraction may reflect a bias to organize information in terms of relational patterns, where certain kinds of patterns (such as those holding between identical elements) are

memorable simply because of their perceptual salience. The purpose of the present experiments was to investigate the extent to which learners acquire knowledge of grammatical sequential dependencies, repetition structure, or both.

The participants in Experiment 1 were trained on a grammar with no repetitions. Half were tested with their training vocabulary, to assess learning of the grammar's sequential structure. The other half were tested with a new vocabulary. On the sequential dependency view of abstraction, repeating elements should have no special status, and therefore, positive transfer after exposure to a grammar with no repetitions would be evidence that learners are abstracting grammatical sequential dependencies. Failure to find transfer would lend support to the notion that learners are not abstracting grammatical sequential dependencies per se but, rather, are performing abstraction on the basis of repetition patterns (Brooks & Vokey, 1991; Mathews & Roussel, 1997). Experiment 1's grammar was modified for Experiment 2, so that a subset of the strings contained a repeating element. This modification was used to further investigate the effect that repetition structure might have on overall learning.

EXPERIMENT 1

Method

Participants

The participants were 44 undergraduates at New Mexico State University, fulfilling requirements for an introductory psychology course. Ten participants were excluded for responding *no* to every item during transfer and two for responding *yes* to every item on the letter-same test.[2] Thirty-two remained.

Materials

Acquisition materials. Thirty-eight letter strings, ranging in length from three to nine letters, were generated from the grammar shown in Figure 3. These were used as training exemplars during acquisition (see the Appendix). The strings were instantiated in one
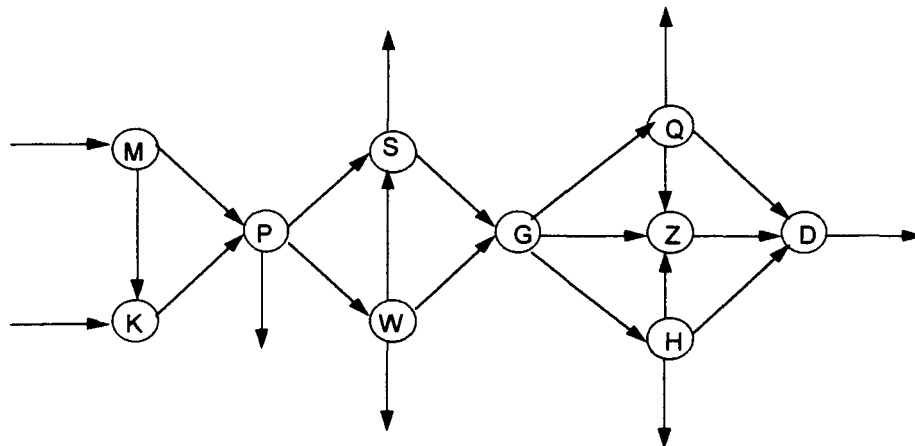
Figure 3. Grammar used in Experiment 1. A Strings were also instantiated in the letters JCXRTBNFVL.

of two letter sets (MKPSWGQZHD and JCXRTBNFVL). Strings were chosen to represent all paths through the grammar.

Test materials. The 34 remaining grammatical letter strings were used during test. Thirty-four ungrammatical strings were constructed by inserting one illegal pairwise transition in an otherwise grammatical test string. For example, the transition PD is illegal in the grammar, so the grammatical string MKPW was made illegal by inserting PD to obtain MKPD. Illegal transitions were inserted in three locations (beginning, middle, or end transitions) in order to test sensitivity to violations across string position. Altogether, there were 68 test items.

Strings were chosen in order to preserve as closely as possible the proportions of strings of similar lengths in the study items. In previous studies, participants had shown considerable learning of string-initial and string-final information (Reber & Allen, 1978; Reber & Lewis, 1977). Therefore, strings never began or ended with illegal letters.

Procedure

The participants were seated in front of PC-compatible computers with standard keyboards. The experiment was entirely automated, but the participants were encouraged to ask for help at any time.

Acquisition phase (typing task). The participants were in one of four groups during acquisition. They were trained on letter strings in the same or in a different letter set than they would first encounter during test and were further grouped according to the letter set they received during training. Letter strings were displayed on a computer screen for an amount of time proportional to the number of letters in the string. For example, a five-letter string was displayed for 1,000 msec (200 msec for each letter). The participants were instructed to study the string and then type it at the keyboard once it disappeared from view. If the participant typed the string correctly, the next string was presented, but in the case of an error, the string reappeared, and the participant had to type it again. This procedure continued until the participant had typed the entire string correctly. The 38 acquisition strings were presented in a different random order in each of four blocks of trials. The participants received feedback on their error rate at the end of each block.

Test. Strings were presented one at a time on the computer screen. The length of presentation was determined in the same manner as during acquisition. Instructions were identical for both the letter-same and the transfer conditions. The participants were informed that the strings used in the typing task were determined by complex rules allowing certain letters to follow others. They were warned that they would have to make judgments on strings in a new

letter set in one of the conditions but that grammatical strings would, nevertheless, conform to the rules used during training. The participants were then asked to discriminate new strings that were valid instances of the rules from those violating the rules. The participants made their choice by pressing the "y" key on the keyboard if they thought the string was grammatical and the "n" key if the string was ungrammatical.

Design

The participants received four test blocks, two letter same and two letter different. Each block contained the 68 test items in random order. Two blocks of each type were used in order to test two predictions consistent with the Dienes et al. (1999) model: (1) Letter-same performance should decrease slightly if the model continues to learn during the test, and (2) transfer performance should improve as the mapping to new vocabulary is learned. In addition, vocabulary (letter same vs. letter different) was counterbalanced so that half of the participants were tested in their training vocabulary before being tested for transfer. The other half were transferred to new vocabulary first. The analyses took the form of 2 × 2 × 2 × 2 (test order × letter set × block × grammaticality [grammatical vs. ungrammatical]) mixed design analyses of variance (ANOVAs). Separate analyses were conducted for letter-same and transfer conditions. Analyses revealed no effects of, or interactions with, test order or letter set. Mean endorsement rates to grammatical and ungrammatical strings were used as the dependent measure.

Results and Discussion

Overall String Analyses

Mean endorsement rates on the grammaticality test are shown in Table 1.

Letter-same performance. The participants, tested in their training vocabulary, were more likely to endorse

Table 1
Mean Endorsement Rates Obtained for Overall Analyses in Experiment 1 After Exposure to a Grammar With No Repetitions

Status of Grammaticality	Vocabulary			
	Letter Same		Transfer	
	M	SD	M	SD
Grammatical	.71	.11	.57	.11
Ungrammatical	.50	.11	.54	.10

grammatical than ungrammatical strings, resulting in a main effect of grammaticality [$F(1,28) = 49.58$, $MS_e = 0.03$, $p \leq .001$]. There was also a significant grammaticality × block interaction [$F(1,28) = 4.01$, $MS_e = 0.006$, $p = .05$]. Learners showed a greater mean difference in endorsement rates between grammatical and ungrammatical strings in the first ($M_{diff} = .25$) compared with the second ($M_{diff} = .19$) test block. This interaction is predicted by the Dienes et al. (1999) model when learning continues (i.e., the weights are changed) during the test.

Transfer performance. Despite robust letter-same performance, learners showed no grammaticality effect during transfer [$F(1,28) = 2.92$, $MS_e = 0.01$, $p = .10$]. The test was sensitive enough to detect an effect as small as .041 (where the actual mean difference in performance on grammatical and ungrammatical strings was .03 ± .04, as given by the 95% confidence interval). There were no other main effects or interactions ($Fs < 1$).

Fine-Grained String Analyses

Several studies have shown that learners acquire considerable knowledge of string-initial information in their training vocabulary (Reber & Allen, 1978; Reber & Lewis, 1977). Perhaps the failure to obtain transfer was due to extensive learning of initial dependencies at the expense of those occurring further along in sequence. We investigated this possibility by conducting analyses of participants' ability to detect violations in specific positions in sequence. Mean endorsement rates to violations by position are shown in Table 2. Violations in beginning position were those occurring in the first two transitions ($n = 8$), violations in end position were those occurring in the last two transitions ($n = 14$), and violations in middle position were those occurring in between ($n = 10$). By these definitions, violations in strings three letters in length were ambiguous (violations occurred in both beginning and end positions, $n = 2$). Responses to strings three letters in length were not included in the means reported in the table but are discussed below.

Letter-same performance. As is shown in Table 2, the participants exhibited considerable knowledge of the grammar's sequential dependencies. Learners tested in their training vocabulary detected violations across all three string positions ($Fs \geq 5.45$, $ps \leq .027$). There was

Table 2
Mean Endorsement Rates (With Standard Deviations) Obtained for Position Analyses in Experiment 1 After Exposure to a Grammar With No Repetitions

Status of Grammaticality	Vocabulary			
	Letter Same		Transfer	
	M	SD	M	SD
Grammatical	.70	.11	.56	.10
Ungrammatical position of violation				
Beginning	.36	.22*	.41	.19*
Middle	.63	.14*	.60	.14
End	.48	.18*	.54	.14

*$p \leq .05$.

also a significant grammaticality × block interaction for violations in the beginning position [$F(1,28) = 9.89$, $MS_e = 0.02$, $p = .004$]. Learners showed greater mean differences in endorsement rates to grammatical and ungrammatical strings in the first ($M_{diff} = .41$) than in the second ($M_{diff} = .26$) test block. Analyses of the means for strings three letters in length also resulted in a significant grammaticality effect [$F(1,28) = 16.12$, $MS_e = 0.06$, $p \leq .001$], with mean endorsement rates for grammatical and ungrammatical strings equal to .68 and .51, respectively.

Transfer performance. Although endorsement rates were not significantly higher for grammatical than for ungrammatical strings when violations occurred in the string-middle or string-end position, or for strings three letters in length, there was a significant effect for violations in the first transition in sequence [$F(1,28) = 13.86$, $MS_e = 0.05$, $p \leq .001$]. This grammaticality effect occurred regardless of whether learners participated in transfer immediately after acquisition or after initial testing in their training vocabulary. Ninety-five percent confidence intervals were computed as an additional check on sensitivity. The tests were sensitive enough to detect effects as small as .051 for violations in the string-middle position ($M_{diff} = -0.04 \pm .05$), .041 ($M_{diff} = 0.02 \pm .04$) for violations at string end, and .071 for strings three letters in length ($M_{diff} = .03 \pm .07$).

Experiment 1 tested the hypothesis that, if learners abstract the grammar's sequential dependencies, they should be able to transfer after exposure to a grammar without repetitions. Learners acquired robust knowledge of grammatical sequential dependencies, as was evidenced by letter-same performance. However, such knowledge played only a limited role in transfer. Those participants who showed substantial learning of sequential dependencies in the letter-same test were unable to transfer such knowledge beyond the first transition in the sequence.

Although learners appeared to be abstracting the sequential structure of the first transition, such a result cannot be taken as evidence that they have abstracted the entire grammar. What then do learners abstract? According to Mathews and Roussel (1997) and Brooks and Vo-key (1991), learners abstract patterns of repetition across sequential structure. We investigated this hypothesis by modifying Grammar 1 so that one third of its strings contained a repeating element. Learners were then tested to see whether they would distinguish legal from illegal repetitions.

EXPERIMENT 2

Method

Participants

The participants were 42 undergraduates at New Mexico State University, fulfilling requirements for an introductory psychology course. Two were excluded for responding *no* to every item on the transfer test, resulting in 40 remaining participants.

Materials

The grammar used in Experiment 2 is shown in Figure 4. This grammar is identical to that shown in Figure 1, with the exception
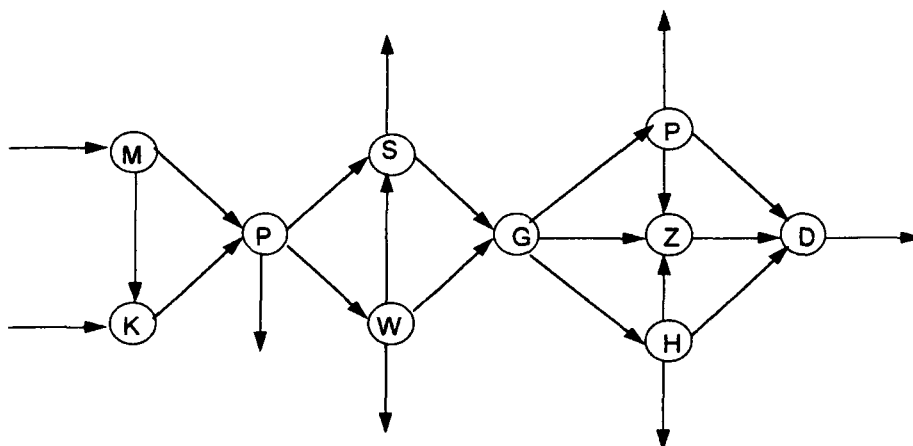
Figure 4. Grammar used in Experiment 2. This grammar was modified so that one third of the strings contained a repeating letter.

of one repeating element (e.g., the letter Q in Figure 3 was replaced by P). The end result was that one third of the strings produced by the grammar contained one repetition. The other strings contained no repetitions and were used as a control for ensuring that learners were indeed acquiring the sequential structure of the grammar. The strings were instantiated in the same letter sets as those used in Experiment 1.

Acquisition materials. Thirty-eight letter strings, ranging in length from three to nine letters, were generated from the grammar shown in Figure 4 (see the Appendix). Each string was presented four times.

Test materials. The 34 remaining grammatical letter strings were used during test. As in Experiment 1, 34 additional ungrammatical strings were constructed by inserting one illegal pairwise transition in an otherwise grammatical test string. Strings with ungrammatical repetition patterns contained one illegal transition and one illegal repetition (repetitions in grammatical strings were always separated by two or three intervening letters, whereas repetitions in ungrammatical strings were separated by only one). For example, the repetition pattern involving S_S in the string MPSWSGH is illegal, but every transition except SW is legal. Each of the strings were presented twice during test.

Procedure

Learners participated in the typing task followed by one grammaticality test (consisting of two blocks), either in their acquisition or in the transfer vocabulary. Otherwise, the procedure was identical to that used in Experiment 1.

Design

The design was the same as that in the previous experiment, with two exceptions. The participants received one grammaticality test only (letter same vs. letter different). Each participant was tested on two string types (with and without repetitions). The analyses took the form of 2 × 2 × 2 × 2 (letter set × block × string type × grammaticality) mixed design ANOVAs. Separate analyses were conducted for letter-same and transfer conditions. Analyses revealed no effects of, or interactions with, letter set. Mean endorsement rates to grammatical and ungrammatical strings were used as the dependent measure. The string KPSPW, which differed from the other strings with illegal repetitions in that it contained the same repeating letter as the grammatical items, was included by mistake. This string was not included in the analyses.

Results and Discussion

Overall Analyses

Mean endorsement rates on the grammaticality test are shown in Table 3.

Letter-same performance. There were main effects of string type [$F(1,18) = 40.45, MS_e = 0.03$] and grammaticality [$F(1,18) = 72.43, MS_e = 0.05, ps < .001$]. The interaction of these variables reflected a greater grammaticality effect for strings with than for those without repetitions [$F(1,18) = 9.82, MS_e = 0.04, p = .006$]. An ANOVA on strings with no repetitions resulted in a significant grammaticality effect [$F(1,18) = 39.70, MS_e = 0.02, p < .001$].

Transfer performance. There was a main effect of grammaticality [$F(1,18) = 4.43, MS_e = 0.04, p = .049$], but of particular note was the interaction of grammaticality with string type [$F(1,18) = 6.27, MS_e = 0.03, p = .022$]. Strings with repetitions clearly had a greater grammaticality effect than strings without them. An ANOVA on strings with repetitions resulted in a significant effect of grammaticality [$F(1,18) = 8.24, MS_e = 0.05, p = .01$]. An ANOVA on strings with no repetitions resulted in a null grammaticality effect ($F < 1$). The test was sensitive enough to detect an effect as small as .051 ($M_{diff} = .02 \pm .05$, as given by the 95% confidence interval).

Table 3
Mean Endorsement Rates (With Standard Deviations)
Obtained in Overall Analyses for Experiment 2

Condition	With Repetitions		Without Repetitions	
	M	SD	M	SD
Letter same				
Grammatical	.65	.16	.73	.10
Ungrammatical	.26	.20	.53	.17
Transfer				
Grammatical	.57	.23	.55	.18
Ungrammatical	.42	.21	.53	.16

Table 4
Mean Endorsement Rates (With Standard Deviations)
Obtained for Position Analyses in Experiment 2

| | String Type | | | |
| | With Repetitions | | Without Repetitions | |
Condition	M	SD	M	SD
Letter same				
Grammatical	.65	.16	.71	.10
Ungrammatical position of violation				
Beginning	.19	.16*	.46	.33*
Middle	.37	.29*	.65	.23
End	.24	.17*	.48	.18*
Transfer				
Grammatical	.57	.23	.54	.19
Ungrammatical position of violation				
Beginning	.31	.33*	.44	.31
Middle	.42	.26*	.52	.20
End	.52	.20	.53	.18

$*p \leq .05$.

Fine-Grained Analyses

Analyses specific to the position of transition violations were used to investigate, first, whether the participants were learning transitions across all positions in the grammar and, second, the extent to which this knowledge was transferred to new vocabulary. Beginning, middle, and end violations were defined in the same way as those in Experiment 1. For strings with no repetitions, 4 had string-beginning violations, 6 had string-middle violations, and 10 had string-end violations. For strings with repetitions, 4 had string-beginning violations, 3 had string-middle violations, and 4 had string-end violations. Mean endorsement rates to violations by position are shown in Table 4.

Letter-same performance. Both grammaticality and string type showed main effects across all three positions ($Fs \geq 12.117, ps \leq .003$). Interactions of these variables demonstrated greater grammaticality effects for repetitions ($Fs \geq 5.97, ps \leq .025$). Nevertheless, the participants exhibited robust learning of the grammars' sequential dependencies. ANOVAs on strings without repetitions showed significant detection of violations in the beginning and end positions and also in strings with only three letters ($Fs \geq 10.38, ps \leq .005$; mean endorsement rates were .89 and .56 for grammatical and ungrammatical strings).

Transfer performance. This group showed a different pattern of effects. There were main effects of grammaticality in the beginning and middle positions in sequence, ($Fs \geq 6.82, ps \leq .017$), but of greater interest were interactions of these variables with string type. The interaction for middle position violations was significant [$F(1,18) = 5.01, MS_e = 0.04, p = .038$] and, for beginning position violations, was marginally significant, [$F(1,18) = 3.70, MS_e = 0.07, p = .070$]. ANOVAs on strings with repetitions showed that learners were able to detect illegal repetitions in string-beginning [$F(1,18) = 13.83, MS_e = 0.10, p = .002$] and string-middle [$F(1,18) = 7.24, MS_e = $

$0.07, p = .015$] positions. ANOVAs on strings without repetitions showed that learners were unable to detect violations in any position ($Fs \leq 2.61, ps \geq .123$). Confidence intervals, for performance on strings without repetitions, were computed as an additional check on sensitivity. The tests were sensitive enough to detect effects as small as .131 for violations in string-beginning position ($M_{diff} = .10 \pm .13$), .051 for violations in string-middle positions ($M_{diff} = .01 \pm .05$), .061 for violations at string-end positions ($M_{diff} = .01 \pm .06$), and .141 for strings three letters in length ($M_{diff} = -.06 \pm .14$).

Experiment 2 examined the hypothesis that learners use patterns of repetition structure to discriminate legal from illegal strings during transfer. Grammar 1 was modified so that one third of its strings contained a repeating element. Learners were then tested to see whether they would distinguish legal from illegal repetitions. Learners were also tested on strings with no repetitions. Learners acquired considerable knowledge of sequential dependencies in their acquisition vocabulary, as measured by performance on strings with no repetitions. However, such knowledge was not used during transfer. Learners could not discriminate grammatical and ungrammatical strings in new vocabulary when strings had no repetitions. They were, however, able to discriminate legal from illegal repetition patterns.

It is interesting to note that endorsement rates on letter-same strings with no repetitions were virtually identical in Experiments 1 and 2. Nevertheless, the fact that endorsement rate differences were significantly greater for strings with repetitions than for strings without them suggests that repetition structure is particularly salient for learners. Recall that the learners in Experiment 1 showed a grammaticality effect for violations in the first transition during transfer. They also showed a grammaticality effect for violations in the middles of strings, when tested in their training vocabulary. Neither of these effects was replicated in the present experiment, suggesting that the mere presence of repetition structure may change the nature of learning. That is, learners may capitalize on the perceptual saliency of this structure at the expense of acquiring sequential knowledge.

Note in Table 3 that learners were far better at rejecting letter-same strings with ungrammatical repetitions ($M = .26$) than illegal sequential dependencies ($M = .53$), suggesting that illegal repetitions are particularly noticeable in the letter-same condition. It is important to rule out the possibility that this pattern signals the presence of factors unrelated to transfer. The lower endorsement rates on strings with illegal repetitions may have been made more salient by including two sources of illegality (an illegal transition and an illegal repetition). The fact that illegal repetitions were not as noticeable during transfer, in combination with the fact that learners did not transfer sequential structure, supports this interpretation. Note also that endorsement rates for grammatical strings with rep-

etitions were somewhat lower ($M = .65$) than those for strings without them ($M = .73$). Although this difference reflects greater learning of sequential dependencies than of repetitions, this slight edge was of no use during transfer.

GENERAL DISCUSSION

The purpose of the present experiments was to determine the degree to which learners use knowledge of sequential dependencies, repetition structure, or both to transfer to new vocabulary as a result of artificial grammar learning. According to one view of transfer, learners abstract the grammar's sequential dependencies and then learn a mapping to new vocabulary at test (Altmann et al., 1995; Dienes et al., 1999). Repeating elements have no special status on this view, and so a logical prediction is that learners should still transfer after exposure to a grammar without such elements. Repetition structure is absolutely necessary for transfer, according to the other view, (Brooks & Vokey, 1991; Mathews & Roussel, 1997).

The participants in Experiment 1 were trained on a grammar with no repetitions and were then tested in their training and transfer vocabulary. Learners acquired considerable knowledge of sequential dependencies, as measured by test performance in their training vocabulary. Such knowledge was evident across beginning, middle, and end transitions in the grammar. The grammars used in most other studies have contained repeating elements, and so this finding is important for demonstrating that participants have learned transitions across the entire grammar.[3] However, the same participants who performed well in their training vocabulary showed only limited transfer to new vocabulary, suggesting that knowledge of sequential dependencies factors little into transfer.

Experiment 2 replicated and extended the findings from the first experiment. The participants were trained on a slightly modified grammar in which one third of the strings contained a repeating element. Letter-same performance showed that the participants acquired the sequential structure of the grammar but that such knowledge was not transferred. The ability to distinguish grammatical and ungrammatical strings was demonstrated only on strings with repetition structure.

Repetition structure appears to be salient for younger learners also. Gomez and Gerken (1997, 1998, 1999) trained 11-month-old infants on grammars with and without repetitions. Strings were presented as sequences of CVC syllables, using a head-turn preference procedure in which auditory stimuli were presented to the infant's left or right side (Kemler Nelson et al., 1995). The dependent measure was the amount of time infants oriented toward the source of sound. Infants were exposed to strings from one of two grammars with no repetitions. When tested in their training vocabulary, infants showed a listening preference for strings from their training gram-

mar, as compared with strings from the other grammar, but not when transferred to new vocabulary (Gomez & Gerken, 1998). Infants do transfer, however, after exposure to grammars with repetitions. Gomez and Gerken (1997, 1999) trained 11-month-olds on one of two grammars with different repetition patterns and then tested them in new vocabulary. Infants showed a listening preference for strings from their training grammar versus strings from the other grammar, regardless of which grammar they were trained on. These results are remarkable, given that infants accumulated an average of only 2 min of exposure to the grammar during training.

In summary, learners acquire the grammar's sequential dependencies but do not appear to abstract them; rather, they abstract repetition structure. This finding is consistent both with Brooks and Vokey (1991) and with Mathews and Roussel (1997). Such findings do not imply that a simple recurrent network (e.g., of the type proposed by Dienes et al., 1999) is inadequate for explaining *any* abstraction involving sequential structure but, rather, that certain architectures do not accurately model the transfer observed in artificial grammar learning thus far. Connectionist networks exhibit a number of properties desirable for modeling implicit learning (Cleeremans, 1997), and so one approach to devising an appropriate model might be to make repetitions of elements an explicit part of the architecture. Such a modification would mean that the same architecture could be used to predict letter-same performance on grammars with and without repetitions, as well as for transfer.

Given the conclusions reached in this paper, several competing explanations should be considered. One has to do with sufficient checks on sensitivity. Despite the interaction between repetition and no-repetition conditions in Experiment 2, the conclusions drawn from these experiments are largely based on null results. Confidence intervals were reported for null effects, to demonstrate that actual performance was well below that detectable by the statistical test. However, another way to test sensitivity is to assume that transfer will always be some fraction of letter-same performance. In the repetition condition (Experiment 2), the letter-same effect (mean endorsement rate for grammatical vs. ungrammatical strings) was .39, whereas the transfer effect (.15) was 39% of letter-same performance. This size effect is, in fact, similar to those reported for previous studies (see, e.g., Altmann et al., 1995; Dienes & Altmann, 1997). In contrast, the letter-same effect in the no-repetition condition was .20. If we assume that transfer can be expected to be about $.39 \times .20 = .08$, we might ask whether our experiments could have detected this size effect. Given the obtained confidence interval reported in Experiment 2 for transfer performance on strings with no repetitions ($\pm .05$), our tests were sensitive enough to detect the expected effect of .08, and yet the reported effect was .02. Similarly, in Experiment 1, the confidence interval

($\pm .04$) would also have allowed us to detect the .08 effect, but the obtained effect was only .03.

A second, competing explanation has to do with the relative difficulty of abstracting repetition versus sequential structure. What if abstracting sequential dependencies is not impossible but simply harder than abstracting repetition structure? Then, we might ask whether more exposure would have resulted in greater success. This explanation is unlikely. The participants in our experiments were given four times the exposure obtained by learners in previous studies (cf. Altmann et al., 1995; Tunney & Altmann, 1999) and were still unable to transfer sequential structure.

Third, could the advantage for repetitions have been an artifact of the typing task? That is, were learners more likely to detect repetitions because they involved identical key strokes? This too is unlikely. Tunney and Altmann (1999) tested learners to see whether they would detect multiple sequential violations in strings with legal repetition structure. Rather than typing strings, however, learners studied them. Learners were no more successful at detecting violations of sequential structure during transfer as a result of these training conditions than they were in our experiments.

In conclusion, transfer has been particularly intriguing, because it suggests that learning is not tied to orderings of specific units but, rather, to orderings of abstract units. Missing from this account, however, has been a more precise understanding of what learners actually abstract. The present experiments take us a step closer to understanding transfer by demonstrating that learners abstract repetitions rather than sequential dependencies. Such learning may reflect a bias to organize information in terms of relational patterns, where certain kinds of patterns (such as those holding between identical elements in sequence) are memorable because of their perceptual salience. For example, recognizing MPSGP and JXRBX as instances of the pattern _x_ _x entails noting the identity relation between the second and the fifth letters in the sequence.

Pattern-based abstraction is important because it signals a system capable of representing relations among abstract variables (cf. Marcus et al., 1999), but what are the implications of such learning? Relational patterns abound in natural language. An example found in English is the pattern *noun verb noun*. Unlike the repetition structure found in finite-state grammars, however, this relational pattern is not cued by the physical stimulus but, rather, by abstract categories. That is, abstraction of the pattern x _ x from the string PSP involves determining that the first and third elements are identical, whereas abstraction of noun verb noun from the sentence *dogs chase cats*, involves identifying the first and third elements as members of the category noun and the second as a verb. The latter determination cannot be made solely on the basis of identity relations among elements in the physical stimulus. *Dogs chase cats* and *John writes papers* share the same category-based structure, despite the physical dis-

similarities among category members. Thus, pattern-based and category-based abstraction appear to differ in fundamental ways. Additional research will be needed to determine the extent to which these two forms of abstraction reflect the same or different learning mechanisms.

REFERENCES

ABRAMS, M., & REBER, A. S. (1988). Implicit learning: Robustness in the face of psychiatric disorders. *Journal of Psycholinguistic Research*, 17, 425-439.

ALLEN, R., & REBER, A. S. (1980). Very long term memory for tacit knowledge. *Cognition*, 8, 175-185.

ALTMANN, G. T. M., DIENES, Z., & GOODE, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 21, 899-912.

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. (1997). Principles for implicit learning. In D. C. Berry (Ed.), *How implicit is implicit learning?* (pp. 195-234). New York: Oxford University Press.

DIENES, Z., & ALTMANN, G. A. (1997). Transfer of implicit knowledge across domains: How implicit and how abstract? In D. C. Berry (Ed.), *How implicit is implicit learning?* (pp. 107-123). New York: Oxford University Press.

DIENES, Z., ALTMANN, G. T. M., & GAO, S.-J. (1999). Mapping across domains without feedback: A neural network model of transfer of implicit knowledge. *Cognitive Science*, 23, 53-82.

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

ELMAN, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.

GOMEZ, R. L. (1997). Transfer and complexity in artificial grammar learning. *Cognitive Psychology*, 33, 154-207.

GOMEZ, R. L., & GERKEN, L. A. (1996, November). *Artificial grammar learning in one-year-olds: Evidence for generalization to new structure.* Paper presented at the 21st Annual Boston University Conference on Language Development.

GOMEZ, R. L., & GERKEN, L. A. (1997, November). *Abstraction in infant artificial grammar learning.* Paper presented at the 22nd Annual Boston University Conference on Language Development.

GOMEZ, R. L., & GERKEN, L. A. (1998, April). *The basis of abstraction in infant artificial grammar learning.* Poster presented at the 13th Biennial International Conference on Infant Studies, Atlanta.

GOMEZ, R. L., & GERKEN, L. A. (1999). Artificial grammar learning by one-year-olds leads to specific and abstract knowledge. *Cognition*, 70, 109-135.

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

KEMLER NELSON, D. G., JUSCZYK, P. W., MANDEL, D. R., MYERS, J., TURK, A., & GERKEN, L. A. (1995). The head-turn preference procedure for testing auditory perception. *Infant Behavior & Development*, 18, 111-116.

KNOWLTON, B. J., & SQUIRE, L. R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 20, 79-91.

KNOWLTON, B. J., & SQUIRE, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 22, 169-181.

MARCUS, G. F., VIJAYAN, S., BANDI RAO, S., & VISHTON, P. M. (1999). Rule learning by seven-month-old infants. *Science*, 283, 77-80.

MATHEWS, R. C., BUSS, R., STANLEY, W. B., BLANCHARD-FIELDS, F., CHO, J. R., & DRUHAN, B. (1989). Role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 15, 1083-1100.

MATHEWS, R. C., & ROUSSEL, L. G. (1997). Abstractness of implicit knowledge: A cognitive evolutionary perspective. In D. C. Berry (Ed.), *How implicit is implicit learning?* (pp. 13-47). New York: Oxford University Press.

MEULEMANS, T., & VAN DER LINDEN, M. (1997a). Associative chunk strength in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 23, 1007-1028.

MEULEMANS, T., & VAN DER LINDEN, M. (1997b). Does the artificial grammar learning paradigm involve acquisition of complex information? *Psychologica Belgica*, 37, 69-88.

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 & Verbal Behavior*, 77, 317-327.

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

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

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

SHANKS, D. R., JOHNSTONE, T., & STAGGS, L. (1997). Abstraction processes in artificial grammar learning. *Quarterly Journal of Experimental Psychology*, 50, 216-252.

TUNNEY, R. J., & ALTMANN, G. T. M. (1999). The transfer effect in artificial grammar learning: Re-appraising the evidence on the transfer of sequential dependencies. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 25, 1322-1333.

VOKEY, J. R., & HIGHAM, P. A. (1997, June). *Abstract analogies and implicit abstraction of structure*. Paper presented at the First Conference of the Association for the Scientific Study of Consciousness, Claremont, CA.

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: Learning, Memory, & Cognition*, 19, 227-248.

WHITTLESEA, B. W. A., & WRIGHT, R. L. (1997). Implicit (and explicit) learning: Acting adaptively without knowing the consequences. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 23, 181-200.

NOTES

1. The model proposed by Mathews and Roussel (1997) deals only with immediate repeats, but it is a simple matter to generalize the model to repetitions spanning any number of intervening elements.

2. It is likely that the large number of participants who followed the strategy of making a single response during transfer was due to the absence of repetition structure. Without repetitions, the participants had no apparent means for making discriminations. Some participants made an attempt to guess, but others chose the same response to each test item. The instructions were identical to those used in Experiment 2, where only 2 participants chose the same response for each test item. The difference was that the participants in Experiment 2 were able to use repetition structure to make discriminations.

3. The only study investigating learning of a grammar with no repetitions was reported in Whittlesea and Wright (1997), but learners were exposed to the same set of grammatical strings during training and test.

APPENDIX
Training and Test Items Used in Experiments 1 and 2

Training Items

Experiments 1-2	MKPWGHZD	MKPSGZD	MKPWSGHZD	KPSGHZD
MPW	MPSGH	MPWS	KPS	KPWS
MPWGHD	MPSGHZD	MPWSGHD	KPWGHD	KPWSGHD
MPWGZD	MKPS	MPWSGZD	KPWGZD	KPWSGZD
MKPWGH	MKPSGHD	MKPWSGH	KPSGH	MKP

Experiment 1	MPSGQD	MKPWGQD	MKPWSGQ	KPWGQZD
MPWGQ	MPWSGQ	MKPSGQD	MKPWSGQD	KPSGQD
MPWGQZD	MPWSGQZD	MKPSGQZD	KPWGQ	KPWSGQZD

Experiment 2	MPSGPD	MKPWGPD	MKPWSGP	KPWGPZD
MPWGP	MPWSGP	MKPSGPD	MKPWSGPD	KPSGPD
MPWGPZD	MPWSGPZD	MKPSGPZD	KPWGP	KPWSGPZD

Note—Training items were instantiated in letter sets MKPSWGQZHD or JCXRTBNFVL.

Test Items

Experiments 1-2

Legal	MKPW	MPSGZD	MKPWS	KPW
	MKPWGHD	MKPSGH	MKPWSGHD	KPSGHD
	MKPWGZD	MKPSGHZD	MKPWSGZD	KPSGZD
MPWGH	MPS	MPWSGH	KPWGH	KPWSGH
MPWGHZD	MPSGHD	MPWSGHZD	KPWGHZD	KPWSGHZD

Illegal	MKPD	KMPSGH	MKPWD	KPH
	KMPWGHD	MKPGZD	KMPWSGHD	KPSWGH
	MKPSGZH	MKPSGZH	MKPWSGZH	KPSGZH
MPWGD	MPD	MPSWGH	KPWGZH	KMPWGH
MPSWGZD	MPSGZH	MPSWGHZD	KPWGZHD	KPSWGHZD

APPENDIX (Continued)

Experiment 1				
Legal	MPWGQD	MPWSGQD	MKPSGQ	KPSGQ
	MPSGQ	MKPWGQ	MKPWSGQZD	KPSGQZD
	MPSGQZD	MKPWGQZD	KPWGQ	KPWSGQD
Illegal	K̲MPWGQ	K̲MPWSGQ	MKPS̲GD	K̲MPWS
	MPS̲QD	MKP̲GQD	MKPS̲WGQZD	KPS̲WGQD
	MPSGQZ̲H	MKPWGQZ̲H	KPWG̲D	K̲MPWSGQ
Experiment 2				
Legal	MPWGPD	MPWSGPD	MKPSGP	KPSGP
	MPSGP	MKPWGP	MKPWSGPZD	KPSGPZD
	MPSGPZD	MKPWGPZD	KPWGPD	KPWSGPD
Illegal	K̲MKPWS	MPS̲WSGH	MK̲MPWS	KPS̲PW
	MPWG̲W	MKPWG̲W	MKPWSGHZ̲H	KPSGHZ̲H
	MK̲MPWGH	MKPS̲WSGH	K̲MKPS	KPWS̲WGH

Note—Test items were instantiated in letter sets MKPSWGQZHD or JCXRTBNFVL.