Research Article

Statistical Learning Within and Between Modalities

Pitting Abstract Against Stimulus-Specific Representations

Christopher M. Conway[1] and Morten H. Christiansen[2]

[1]*Indiana University and* [2]*Cornell University*

ABSTRACT—*When learners encode sequential patterns and generalize their knowledge to novel instances, are they relying on abstract or stimulus-specific representations? Research on artificial grammar learning (AGL) has shown transfer of learning from one stimulus set to another, and such findings have encouraged the view that statistical learning is mediated by abstract representations that are independent of the sense modality or perceptual features of the stimuli. Using a novel modification of the standard AGL paradigm, we obtained data to the contrary. These experiments pitted abstract processing against stimulus-specific learning. The findings show that statistical learning results in knowledge that is stimulus-specific rather than abstract. They show furthermore that learning can proceed in parallel for multiple input streams along separate perceptual dimensions or sense modalities. We conclude that learning sequential structure and generalizing to novel stimuli inherently involve learning mechanisms that are closely tied to the perceptual characteristics of the input.*

A core debate in the psychological sciences concerns the extent to which acquired knowledge consists of modality-dependent versus abstract representations. Traditional information-processing approaches to cognition have emphasized the operation of amodal symbol systems (Fodor, 1975; Pylyshyn, 1984), whereas more recently, embodiment and similar theories have proposed instead that cognition is grounded in modality-specific sensorimotor mechanisms (Barsalou, Simmons, Barbey, & Wilson, 2003; Glenberg, 1997). This debate has been especially intense in the area of implicit statistical learning of artificial grammars.[1] In his early work, A.S. Reber (1967, 1969) demonstrated implicit learning in participants who were exposed to letter strings generated from an artificial grammar. The letter strings obeyed the overall rule structure of the grammar, being constrained in terms of which letters could follow which other letters. Participants not only showed evidence of learning this structure implicitly, but also could apparently transfer their knowledge of the legal regularities from one letter vocabulary (e.g., *M*, *R*, *T*, *V*, *X*) to another (e.g., *N*, *P*, *S*, *W*, *Z*) as long as the same underlying grammar was used for both (A.S. Reber, 1969). This effect has been replicated many times, with transfer being demonstrated not just across letter sets (e.g., Brooks & Vokey, 1991; Mathews et al., 1989; Shanks, Johnstone, & Staggs, 1997), but also across sense modalities (Altmann, Dienes, & Goode, 1995; Manza & Reber, 1997; Tunney & Altmann, 2001).

Transfer effects in artificial grammar learning (AGL) are usually explained by proposing that the learning is based on abstract knowledge, that is, knowledge not directly tied to the surface features or sensory instantiation of the stimuli (Altmann et al., 1995; Pena, Bonatti, Nespor, & Mehler, 2002; A.S. Reber, 1989; Shanks et al., 1997). For instance, the human cognitive system might encode patterns among stimuli in terms of "abstract algebra-like rules" that encode relationships among amodal variables (Marcus, Vijayan, Rao, & Vishton, 1999, p. 79). Such a proposal emphasizes the learning of structural relations among items and deemphasizes the acquisition of information pertaining to specific features of the stimulus elements. Alternatively, participants may learn the statistical structure of the input sequences using associative mechanisms that are sensitive to modality- or stimulus-specific features (e.g., Chang & Knowlton, 2004; Christiansen & Curtin, 1999; Conway

Address correspondence to Christopher M. Conway, Department of Psychology, 1101 E. 10th St., Indiana University, Bloomington, IN 47405, e-mail: cmconway@indiana.edu.

[1]Artificial grammar learning is statistical in the sense that successful test performance can be achieved by encoding something akin to the frequency of chunks of elements (Perruchet & Pacteau, 1990) or by learning the transitional probabilities among consecutive elements (Saffran, Johnson, Aslin, & Newport, 1999).

& Christiansen, 2005; McClelland & Plaut, 1999; Perruchet, Tyler, Galland, & Peereman, 2004).[2]

In this article, we present new evidence from a set of AGL experiments supporting a modality-constrained or embodied view of statistical learning. In three experiments, we used a novel modification of the AGL paradigm to examine the nature of statistical learning within and across modalities. Specifically, in each experiment, we employed two different finite-state grammars in a dual-grammar crossover design in which the grammatical test sequences of one grammar were used as the ungrammatical test sequences for the other grammar. For example, in Experiment 1, participants were exposed to visual sequences from one grammar and auditory sequences from the other grammar. In the test phase, new grammatical sequences from both grammars were presented. Crucially, for each participant, all test items from both grammars were instantiated only visually or only auditorily. In such a crossover design, if participants have learned the abstract rules underlying both grammars, they ought to classify all sequences generated by the grammars, whether they are presented visually or auditorily, as equally grammatical. However, if participants have learned statistical regularities specific to the sense modality in which those regularities were instantiated, they ought to classify a sequence as grammatical only if it is presented in the same sense modality as were the training sequences generated from the same grammar. The data from these experiments follow this latter pattern, suggesting that learners encoded the sequential patterns and generalized their knowledge to novel instances by relying on stimulus-specific, not abstract, representations.

EXPERIMENT 1: MULTIMODAL LEARNING

In Experiment 1, we assessed multimodal learning by presenting participants with auditory tone sequences generated from one grammar and visual color sequences generated from a second grammar. We then tested participants using novel grammatical stimuli from both grammars; half the stimuli were generated from one grammar and the other half were generated from the other grammar, but all sequences were instantiated in only one of the vocabularies (tones or colors). Given our scoring system, in which a classification of a test sequence as grammatical was scored as correct only if the sequence was presented in the sense modality used during training on the corresponding grammar, a null effect of learning (performance level of 50% correct) could mean (a) that participants were unable to adequately learn the statistical regularities or (b) that participants learned the regularities but the knowledge existed in an amodal format that did not retain information regarding sense modality. Accordingly, performance levels significantly above chance would show both that participants learned the statistical regularities from the grammars and that the knowledge was modality-specific. In order to compare dual-grammar learning to performance in the standard AGL paradigm, we employed single-grammar, unimodal learning conditions as a baseline.

Method

Subjects

Forty students (10 in each condition) were recruited from Cornell University undergraduate psychology classes and received extra credit for their participation.

Materials

Two different finite-state grammars, Grammar A and Grammar B (Fig. 1), were used to generate two sets of nonoverlapping stimuli. We used 9 grammatical sequences from each grammar in the training phase and 10 grammatical sequences from each grammar in the test phase; all sequences contained at least three and no more than seven elements. For a given grammar, each letter was mapped onto a color vocabulary (five differently colored squares) or an auditory vocabulary (five pure tones). The five colored squares ranged along a continuum from light blue to green; the colors were chosen such that each was perceptually distinct yet similar enough to the others to make a verbal coding strategy difficult. The five tones had frequencies of 210, 245,

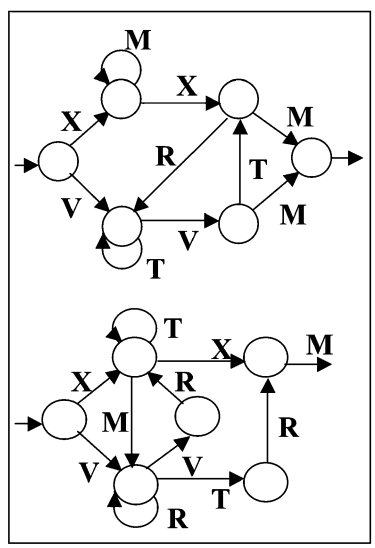

Training Items

Grammar A	Grammar B
VVM	XXM
XMXM	VTRM
XXRVM	VVRXM
VTVTM	XTTXM
XXRVTM	XMRTRM
VTTVTM	VRRTRM
XMMMXM	XMRRTRM
XXRTTVM	VRVRTXM
VTVTRVM	VRRVRXM

Test Items

Grammar A	Grammar B
VTVM	XTXM
VVTM	VRTRM
XMMXM	XTMTRM
VTTVM	VRVRXM
VTTTVM	VVRTXM
VVTRVM	XMVRXM
VTTTVTM	XMVRTXM
XXRTVTM	XMVRTXM
XMXRTVM	VVRTTXM
XMMXRVM	XMRVRXM

Fig. 1. The grammars, training items, and test items used in all three experiments. The diagrams on the left depict Grammar A (top) and Grammar B (bottom). The letters from each grammar were mapped onto colors or tones (Experiment 1), colors or shapes (Experiment 2a), tones or nonwords (Experiment 2b), two different shape sets (Experiment 3a), or two different nonword sets (Experiment 3b).

[2]We distinguish between two related notions of what it means to be abstract (for further discussion, see Dienes, Altmann, & Gao, 1999; Mathews, 1990; Redington & Chater, 1996). Knowledge can be abstract to the extent that it (a) represents common properties among stimuli or (b) is independent of the sense modality or perceptual features of the stimuli. Abstraction in the first sense is generally assumed to be involved in human learning, although it has been hotly debated whether such abstraction occurs via a rule-learning or a statistically based mechanism. The second notion of abstraction has also been a subject of intense debate and is the focus of the current article.

286, 333, and 389 Hz. These frequencies were chosen because they neither conform to standard musical notes nor contain standard musical intervals between them (Conway & Christiansen, 2005). Depending on the experimental condition, the Grammar A sequence *VVM*, for example, might be instantiated as two light-green stimuli followed by a light-blue stimulus or as two 389-Hz tones followed by a 268-Hz tone.

All visual stimuli were presented in a serial format in the center of a computer screen. Auditory stimuli were presented via headphones. Each element (color or tone) of a particular sequence was presented for 500 ms, with 100 ms occurring between elements. A 1,700-ms pause separated each sequence from the next.

Procedure

Participants were randomly assigned to one of four conditions, two experimental and two baseline. Participants in the experimental conditions were exposed to the training sequences from both grammars, with one training set instantiated as colored squares and the other as tones. The assignment of grammars to modalities was counterbalanced across participants. Additionally, within each grammar, the assignment of the letters to particular visual or auditory elements was randomly determined for each participant.

At the beginning of the experiment, participants in the experimental conditions were told that they would hear sequences of auditory stimuli and see sequences of visual stimuli. They were told that it was important to pay attention to the stimuli because afterward they would be tested on what they had observed. The instructions did not explicitly mention the existence of the grammars, nor did they indicate that the sequences followed underlying rules or regularities of any kind. The 18 training sequences (9 from each grammar) were presented randomly, one at a time, in each block, for a total of six blocks. Thus, a total of 108 sequences was presented. Note that because the order of presentation was entirely random, the visual and auditory sequences were completely intermixed with one another. Figure 2 illustrates the stimulus presentation.

In the test phase, participants in the experimental conditions were instructed that the stimuli they had observed were generated according to a complex set of rules that determined the order of the elements within each sequence. Participants were told they would next be exposed to a new set of color or tone sequences. Some of these sequences would conform to the same set of rules as before, whereas the others would be different. Their task was to judge which of the sequences followed the same rules as before and which did not. For the test phase, 20 sequences were used, 10 that were grammatical with respect to one grammar and 10 that were grammatical with respect to the other. For half of the participants, these test sequences were instantiated using the color vocabulary (visual-experimental condition), and for the other participants, the test sequences were instantiated using the tone vocabulary (auditory-experimental condition). For scoring purposes, the test sequences from the grammar that was instantiated in the same sense modality as in the training phase were deemed grammatical, whereas the test sequences from the other grammar were deemed ungrammatical. Thus, a classification judgment was scored as correct if the test sequence was judged as grammatical and its sense modality was the same as that of the training sequences that were generated from the same grammar. Similarly, a classification judgment was also scored as correct if the test sequence was judged as ungrammatical and its sense modality was different from that of the training sequences that were generated from the same grammar. In all other cases, a classification judgment was scored as incorrect.

Participants in the baseline single-grammar conditions followed a similar procedure except that they received training sequences from only one of the grammars, instantiated in just one of the sense modalities, with grammar and modality assignments counterbalanced across participants. The 9 training sequences were presented randomly once per block for six blocks, for a total of 54 presentations. The baseline participants were tested using the same test set as the experimental participants, instantiated with the same vocabulary on which they were trained. Thus, the baseline unimodal conditions

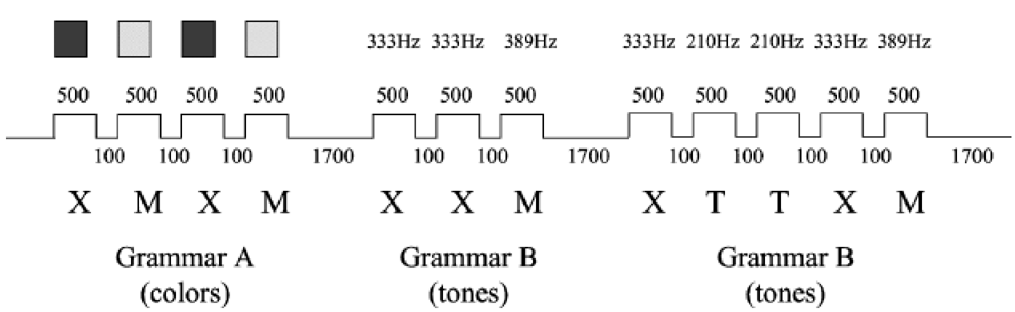

Fig. 2. Sample of stimulus presentation in Experiment 1. Sequences from the two grammars were interleaved randomly. For each participant, one grammar was instantiated with the color vocabulary, and the other grammar was instantiated with the tone vocabulary. Each letter below the time line denotes a particular color or tone, depending on the grammar and vocabulary. The time line indicates the duration of the sequence elements and the intervals between elements, in milliseconds.

(visual-baseline and auditory-baseline conditions) assessed visual and auditory learning with one grammar alone, much as in the standard AGL design.

Results and Discussion

Table 1 reports for each group the mean number and percentage of correct classifications (out of 20), the result of a t test comparing the mean score with chance level, and the p_{rep} value (Killeen, 2005) and effect size, d (Cohen, 1988). Each group's overall performance was better than would be expected by chance. Furthermore, there were no significant differences between the experimental groups and their respective baseline groups: visual-experimental versus visual-baseline, $t(9) < 1$; auditory-experimental versus auditory-baseline, $t(9) = 1.1$, $p = .30$, $p_{rep} = .76$, $d = 0.35$.

These results indicate that participants can simultaneously learn statistical regularities from two input streams generated from two different artificial grammars, each instantiated in a different sense modality. Perhaps surprisingly, performance in the dual-grammar conditions was no worse than performance after single-grammar learning. This lack of a learning decrement suggests that learning of visual statistical structure and learning of auditory statistical structure occur in parallel. Furthermore, these results challenge claims that learning occurs independently of sense modality (e.g., Altmann et al., 1995). If learning had been modality-independent, then test sequences generated by the two grammars would have appeared equally grammatical to the participants, driving performance to chance levels (according to our scoring scheme). Thus, our data suggest that participants' knowledge of the statistical patterns, instead of being amodal or abstract, was stimulus-specific. We next asked whether learners can similarly learn from two different input streams that are within the same sense modality.

EXPERIMENT 2: INTRAMODAL LEARNING ALONG DIFFERENT PERCEPTUAL DIMENSIONS

The purpose of Experiment 2 was to further explore the stimulus-specific nature of implicit statistical learning. Specifically, we assessed whether participants could learn two sets of statistical regularities when they were presented within the same sense modality but instantiated along two different perceptual dimensions. Experiment 2a examined intramodal learning in the visual modality, and Experiment 2b examined auditory learning. For Experiment 2a, one grammar was instantiated with colors, and the other with shapes. For Experiment 2b, one grammar was instantiated with tones, and the other with nonwords.

TABLE 1

Mean Performance and Results of Tests of Significance (Versus Chance) in Experiments 1, 2, and 3

	Experimental conditions (dual-grammar)				Baseline conditions (single-grammar)			
Modality or dimension	Number correct	Percentage correct	$t(9)$	p_{rep}	Number correct	Percentage correct	$t(9)$	p_{rep}
			Experiment 1					
Visual	12.7	63.5	2.76*	.95[a]	12.4	62.0	2.54*	.94[a]
Auditory	14.1	70.5	4.38**	.99[b]	13.1	65.5	3.44**	.97[a]
			Experiment 2a					
Colors	11.9	59.5	2.97*	.96[a]				
Shapes	11.9	59.5	2.31*	.92[b]	13.2	66.0	6.25***	.99[a]
			Experiment 2b					
Tones	13.7	68.5	4.25**	.99[a]				
Nonwords	12.0	60.0	2.58*	.94[a]	12.2	61.0	2.34*	.93[b]
			Experiment 3a					
Shape Set 1	12.0	60.0	2.58*	.93[a]				
Shape Set 2	11.2	56.0	1.65	.85[b]	11.6	58.0	2.95*	.96[a]
			Experiment 3b					
Nonword Set 1	10.9	54.5	1.49	.83[c]				
Nonword Set 2	12.4	62.0	6.47***	.99[a]	13.3	66.5	3.79**	.98[a]

Note. The number correct is out of a possible maximum of 20. All t tests were two-tailed. For the colors and tones conditions in Experiment 2, the baseline conditions were the baseline conditions in Experiment 1 (i.e., visual-baseline and auditory-baseline conditions, respectively). For the Shape Set 1 and Nonword Set 1 conditions in Experiment 3, the baseline conditions were baseline conditions from Experiment 2 (shapes-baseline and nonwords-baseline, respectively).

[a]$d > .8$. [b]$d > .5$. [c]$d > .2$.

*$p < .05$. **$p < .01$. ***$p < .001$.

Method

Subjects

Sixty participants (10 in each condition) were recruited in the same manner as in Experiment 1.

Materials

Experiment 2 incorporated the same grammars and training and test sequences from Experiment 1. Experiment 2a used two visual vocabularies: the same set of colors used in Experiment 1 and a set of five abstract, geometric shapes. The shapes were chosen to be perceptually distinct yet not amenable to a verbal coding strategy. Experiment 2b used two auditory vocabularies: the same set of tones used in Experiment 1 and a set of recordings of five different nonwords (*vot*, *pel*, *dak*, *jic*, and *rud*) spoken by a human speaker (from Gómez, 2002).

Procedure

Participants were randomly assigned to one of six conditions: two for Experiment 2a, two for Experiment 2b, and two new single-grammar baseline conditions. The general procedure was otherwise the same as in Experiment 1. In Experiment 2a, participants were trained on the two visual grammars and then tested on their ability to classify novel sequences instantiated using one of the two vocabularies. In Experiment 2b, participants were trained on both auditory grammars and then tested on novel sequences instantiated using one of the two auditory vocabularies.

The two new baseline conditions provided data for single-grammar performance for the new shape and nonword vocabularies (note that for the analyses of the colors and tones conditions, we used the baseline data from Experiment 1).

Results and Discussion

Table 1 shows that each group's overall performance was better than expected by chance. Furthermore, there were no statistical differences between the experimental groups and their corresponding baseline groups: colors-experimental versus colors-baseline, $t(9) < 1$; shapes-experimental versus shapes-baseline, $t(9) = -1.13, p = .29, p_{\text{rep}} = .77, d = 0.36$; tones-experimental versus tones-baseline, $t(9) < 1$; nonwords-experimental versus nonwords-baseline, $t(9) = -0.178$, $p = .86$, $p_{\text{rep}} = .55$, $d = 0.056$.

The results for Experiments 2a and 2b were similar to those for Experiment 1. Participants were adept at learning two different sets of statistical regularities simultaneously—even when the same sense modality was used for both (shape and color sequences in Experiment 2a, tone and nonword sequences in Experiment 2b). Performance levels were no worse in these dual-grammar conditions than in single-grammar conditions. These results suggest that participants can acquire statistical regularities from two streams of information within the same sense modality, as long as the two streams differ along a major perceptual dimension. A further implication of these results is that participants' knowledge of the underlying statistical structure was stimulus-specific rather than abstract.

EXPERIMENT 3: INTRAMODAL LEARNING ALONG THE SAME PERCEPTUAL DIMENSION

We next looked at dual-grammar learning within the same sense modality when the vocabularies lay along the same perceptual dimension. Experiment 3a incorporated two different sets of visual shapes, and Experiment 3b incorporated two different sets of auditory nonwords.

Method

Sixty participants (10 in each condition) were recruited. Experiment 3 incorporated the same grammars and sequences that were used in Experiments 1 and 2. Experiment 3a employed two visual vocabularies: Shape Sets 1 and 2 (Fig. 3). Shape Set 1 was the same set of shapes used in Experiment 2a; Shape Set 2 was a new set of shapes similar in overall appearance but perceptually distinct from Set 1. Experiment 3b employed the nonword vocabulary used in Experiment 2b (Nonword Set 1), as well as a new nonword set consisting of *tood*, *jeen*, *gens*, *tam*, and *leb* (Nonword Set 2).

Participants were randomly assigned to one of six conditions, two for Experiment 3a, two for Experiment 3b, and two new single-grammar baseline conditions. The general procedure was identical to that for Experiment 2 except that different vocabularies were used. That is, in Experiment 3a, participants were exposed to sequences from both grammars, with one grammar instantiated using Set Shape 1 and the other grammar instantiated using Set Shape 2; subsequently, they were tested on novel sequences generated from both grammars but instantiated using only one of the vocabularies. Similarly, in Experiment 3b, participants received training sequences from both grammars, one grammar instantiated using Nonword Set 1 and

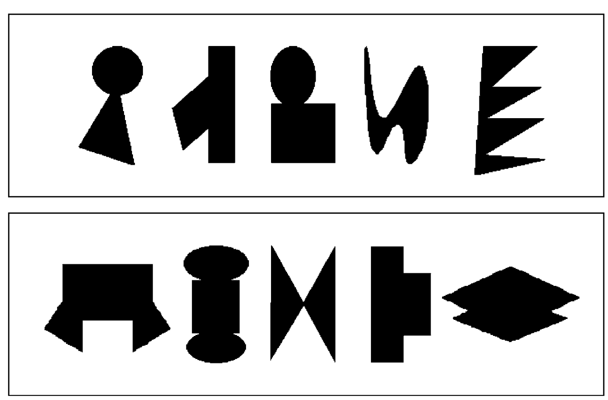

Fig. 3. The visual vocabularies used in Experiment 3a. Shape Set 1 (which was also used in Experiment 2a) is at the top, and Shape Set 2 is at the bottom.

the other generated using Nonword Set 2, and were then tested on sequences generated from both grammars but instantiated using one of the nonword sets only. The two new baseline conditions provided data for single-grammar performance for the new Shape Set 2 and Nonword Set 2 vocabularies.

Results and Discussion

As Table 1 reveals, when exposed to two different statistically governed streams of visual input, each with a distinct vocabulary of shapes, learners on average were able to learn the structure for only one of the streams. This same result was found when learners were exposed to two different nonword auditory streams. Thus, under dual-grammar conditions, learners showed above-chance classification performance for only one of the vocabularies and grammars. As we remarked earlier, chance-level performance could be due to either an inability to learn the underlying regularities or to having acquired these regularities in terms of abstract representations that do not distinguish items on the basis of their perceptual characteristics. Thus, the data from Experiment 3 imply that either (a) intramodal dual-grammar statistical learning did not occur because of perceptual confusion of the stimuli or (b) the knowledge of the two grammars, once learned, was comingled because the input elements were perceptually similar. Either way, traditional theories of AGL that specify abstract representations appear to have difficulty accounting for such low-level, perceptual effects.

OVERALL ANALYSES

To better quantify the differences in learning across the three experiments, we submitted all data to a $4 \times 2 \times 2$ analysis of variance that contrasted condition (multimodal, intramodal–different dimension, intramodal–same dimension, or unimodal baseline), modality (visual or auditory), and grammar (Grammar A or Grammar B). There was a main effect of condition, $F(3, 144) = 2.66, p = .050, p_{\mathrm{rep}} = .92, \eta_p^2 = .053$. There were no main effects of modality or grammar, nor were there any significant interactions (ps $> .05$).

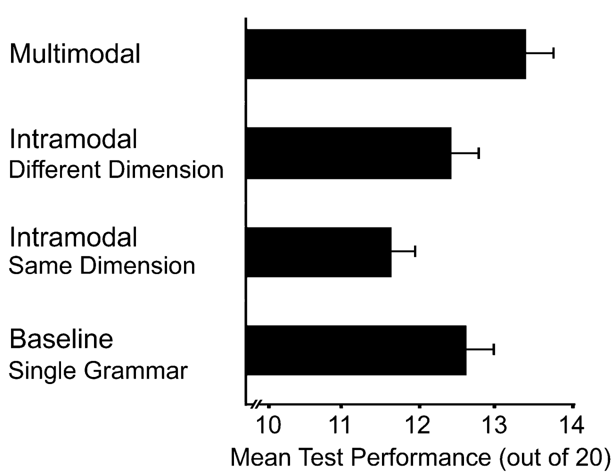

Fig. 4. Mean test performance for all three experiments: multimodal conditions (Experiment 1); intramodal, different-dimension conditions (Experiment 2); intramodal, same-dimension conditions (Experiment 3); and baseline, single-grammar conditions (Experiments 1, 2, and 3).

Figure 4 shows mean test performance collapsed across grammar and modality. Post hoc comparisons revealed that performance in the intramodal, same-dimension condition was significantly lower than performance in both the multimodal ($p = .009, p_{\mathrm{rep}} = .97$) and the baseline ($p = .044, p_{\mathrm{rep}} = .93$) conditions. This outcome confirms that there was a learning decrement for intramodal learning in Experiment 3, when the two grammars were instantiated using vocabularies along the same perceptual dimension.

GENERAL DISCUSSION

In this research, we sought to determine the nature of the acquired knowledge underlying implicit statistical learning. We distinguished between two possibilities. On the one hand, as traditional information-processing approaches suggest, it is possible that learners encode the underlying structure of complex sequential patterns in an abstract (amodal) fashion that does not retain information regarding the perceptual features of the input. On the other hand, embodied cognition theories (Barsalou et al., 2003) suggest that the learner's representations rely on modality-specific sensorimotor systems. Our data support the latter view.

Experiment 1 showed that participants can learn statistical regularities from two artificial grammars when one is presented visually and the other auditorily. Because of our crossover design, the results suggest that learning was modality-specific; otherwise, performance would have been at chance levels. Furthermore, test performance under these multimodal, dual-grammar conditions was identical to performance under unimodal, single-grammar conditions, which suggests that the underlying learning systems operated in parallel and independently of one another. Experiment 2 extended these results, showing that learners can also simultaneously learn regularities from two input streams within the same sense modality—as long as the respective vocabularies differ along a major perceptual dimension. Experiment 3 further showed that learning suffered when the two grammars used vocabularies along the same perceptual dimension; in this case, statistical learning was limited to just one of the two input streams.

These data challenge claims that learning in an AGL task may consist of modality-independent representations (Altmann et al., 1995) or abstract rules (Marcus et al., 1999; A.S. Reber, 1989). Some AGL studies purportedly show transfer effects across modalities, suggesting that the underlying knowledge is abstract and independent of the vocabulary used during training. However, there has been considerable controversy surrounding the transfer data (e.g., Christiansen & Curtin, 1999; Marcus, 1999; Mathews, 1990; McClelland & Plaut, 1999; Redington & Chater, 1996). For example, transfer may be achieved by noticing the presence of low-frequency illegal starting elements in the transfer set (Tunney & Altmann, 1999), rather than by relying on abstract knowledge acquired at training. Or participants may appear to demonstrate transfer if they merely encode certain patterns of repeating elements (e.g., "BD*CCC*B") and then, during the test phase, recognize the same repetition patterns in items with a new vocabulary (e.g., "MT*VVV*M"; Brooks & Vokey, 1991; Redington & Chater, 1996). Thus, it is far from clear that transfer effects reflect the operation of abstract knowledge formed during the learning phase.

In addition to providing evidence for modality-specificity, the data reveal, quite remarkably, that participants are just as adept at learning statistical regularities from two input streams as from one. This points to the possibility of parallel, independent learning mechanisms across and within sense modalities. It has been commonly assumed that statistical learning involves a single, unitary mechanism that operates over all types of input (e.g., Kirkham, Slemmer, & Johnson, 2002). However, our data indicate that this view is inaccurate, or at least incomplete. It is not clear how a single, amodal mechanism could afford simultaneous learning of multiple statistical regularities and keep the stimulus-specific representations independent of one another (Experiments 1 and 2). Previous research has suggested that although there are commonalities in statistical learning across vision, audition, and touch, there also are important modality differences; these findings highlight the possibility of distributed modality-constrained subsystems (Conway & Christiansen, 2005). Such a view of statistical learning resonates with theories of implicit sequence learning (Goschke, Friederici, Kotz, & van Kampen, 2001; Keele, Ivry, Mayr, Hazeltine, & Heuer, 2003), implicit memory (Schacter, Chiu, & Ochsner, 1993), and temporal processing (Mauk & Buonomano, 2004).

Implicit memory research, in particular, may offer insights into the nature of statistical learning. It appears likely that both implicit statistical learning and perceptual priming are supported by something akin to perceptual fluency (Chang & Knowlton, 2004; Kinder, Shanks, Cock, & Tunney, 2003). That is, networks of neurons in modality-specific brain regions show decreased activity when processing items that are the same or similar in overall structure—possibly because of increased processing efficiency for that class of stimuli (P.J. Reber, Stark, & Squire, 1998; Schacter & Badgaiyan, 2001). An explanation of statistical learning in terms of perceptual priming or fluency is consistent with the stimulus-specific learning we observed in the current experiments and may offer the attractive possibility of unifying implicit learning and implicit memory phenomena.

Although the current data point toward modality-specificity, it is possible that learners formed both abstract and stimulus-specific representations, but that the latter were stronger and thus were displayed more readily in the test. Another possibility is that human cognition relies on stimulus-specific representations for some tasks, but abstract learning for others. For example, explicit problem-solving tasks sometimes tap participants' use of abstract principles (Goldstone & Sakamoto, 2003; Reeves & Weisberg, 1994). The ability to learn abstract principles and transfer them to new domains certainly appears to be a hallmark of explicit cognition; it is much less clear, especially in light of the current data, whether it is also a hallmark of implicit learning.[3]

In sum, much of perception and cognition involves the use of multiple sense modalities to implicitly extract structure from temporal or spatiotemporal patterns. The current experiments suggest that the knowledge underlying such implicit statistical learning is closely tied to the sensory and perceptual features of the material, perhaps indicating the involvement of multiple learning subsystems, and challenging traditional theories positing abstract or amodal cognitive processes.

Acknowledgments—This research was conducted as part of the first author's Ph.D. dissertation at Cornell University and was supported in part by Human Frontiers Science Program Grant RGP0177/2001-B to M.H.C. Some of these data were presented at the 27th Annual Conference of the Cognitive Science Society (July 2005 in Stresa, Italy) and the 45th Annual Meeting of the Psychonomic Society (November 2004 in Minneapolis, MN). We wish to thank Rebecca Gómez for providing us with some of the stimuli that were used in these experiments and Thomas Farmer for his help with statistical analyses.

REFERENCES

Altmann, G.T.M., Dienes, Z., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 899–912.

Barsalou, L.W., Simmons, W.K., Barbey, A.K., & Wilson, C.D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*, *7*, 84–91.

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.

Chang, G.Y., & Knowlton, B.J. (2004). Visual feature learning in artificial grammar classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 714–722.

Christiansen, M.H., & Curtin, S. (1999). Transfer of learning: Rule acquisition or statistical learning? *Trends in Cognitive Sciences*, *3*, 289–290.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). New York: Academic Press.

Conway, C.M., & Christiansen, M.H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 24–39.

[3]A further insight comes from the category-learning literature, which posits two systems: an explicit, verbalizable rule-based system and an implicit, procedural-based system (e.g., Maddox & Ashby, 2004). The former is more flexible in that a single verbalizable rule defines the category boundary and thus presumably can be transferred to other domains. The latter system involves learning more complex category boundaries that are nonverbalizable, but are instead tied to specific stimulus-response associations.

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.

Fodor, J.A. (1975). *The language of thought*. New York: Thomas Y. Crowell.

Glenberg, A.M. (1997). What memory is for. *Behavioral and Brain Sciences*, *20*, 1–55.

Goldstone, R.L., & Sakamoto, Y. (2003). The transfer of abstract principles governing complex adaptive systems. *Cognitive Psychology*, *46*, 414–466.

Gómez, R.L. (2002). Variability and detection of invariant structure. *Psychological Science*, *13*, 431–436.

Goschke, T., Friederici, A.D., Kotz, S.A., & van Kampen, A. (2001). Procedural learning in Broca's aphasia: Dissociation between the implicit acquisition of spatio-motor and phoneme sequences. *Journal of Cognitive Neuroscience*, *13*, 370–388.

Keele, S.W., Ivry, R., Mayr, U., Hazeltine, E., & Heuer, H. (2003). The cognitive and neural architecture of sequence representation. *Psychological Review*, *110*, 316–339.

Killeen, P.R. (2005). An alternative to null-hypothesis significance tests. *Psychological Science*, *16*, 345–353.

Kinder, A., Shanks, D.R., Cock, J., & Tunney, R.J. (2003). Recollection, fluency, and the explicit/implicit distinction in artificial grammar learning. *Journal of Experimental Psychology: General*, *132*, 551–565.

Kirkham, N.Z., Slemmer, J.A., & Johnson, S.P. (2002). Visual statistical learning in infancy: Evidence for a domain-general learning mechanism. *Cognition*, *83*, B35–B42.

Maddox, W.T., & Ashby, F.G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioral Processes*, *66*, 309–332.

Manza, L., & Reber, A.S. (1997). Representing artificial grammars: Transfer across stimulus forms and modalities. In D.C. Berry (Ed.), *How implicit is implicit learning?* (pp. 73–106). Oxford, England: Oxford University Press.

Marcus, G.F. (1999). Connectionism: With or without rules? Response to J.L. McClelland and D.C. Plaut. *Trends in Cognitive Sciences*, *3*, 168–170.

Marcus, G.F., Vijayan, S., Rao, S.B., & Vishton, P.M. (1999). Rule learning by seven-month-old infants. *Science*, *283*, 77–79.

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.

Mauk, M.D., & Buonomano, D.V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience*, *27*, 307–340.

McClelland, J.L., & Plaut, D.C. (1999). Does generalization in infant learning implicate abstract algebra-like rules? *Trends in Cognitive Sciences*, *3*, 166–168.

Pena, M., Bonatti, L.L., Nespor, M., & Mehler, J. (2002). Signal-driven computations in speech processing. *Science*, *298*, 604–607.

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

Perruchet, P., Tyler, M.D., Galland, N., & Peereman, R. (2004). Learning nonadjacent dependencies: No need for algebraic-like computations. *Journal of Experimental Psychology: General*, *133*, 573–583.

Pylyshyn, Z.W. (1984). *Computation and cognition: Toward a foundation of cognitive science*. Cambridge, MA: MIT Press.

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

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.

Reber, P.J., Stark, C.E.L., & Squire, L.R. (1998). Cortical areas supporting category learning identified using functional MRI. *Proceedings of the National Academy of Sciences, USA*, *95*, 747–750.

Redington, M., & Chater, N. (1996). Transfer in artificial grammar learning: A reevaluation. *Journal of Experimental Psychology: General*, *125*, 123–138.

Reeves, L.M., & Weisberg, R.W. (1994). The role of content and abstract information in analogical transfer. *Psychological Bulletin*, *115*, 381–400.

Saffran, J.R., Johnson, E.K., Aslin, R.N., & Newport, E.L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70*, 27–52.

Schacter, D.L., & Badgaiyan, R.D. (2001). Neuroimaging of priming: New perspectives on implicit and explicit memory. *Current Directions in Psychological Science*, *10*, 1–4.

Schacter, D.L., Chiu, C.Y.P., & Ochsner, K.N. (1993). Implicit memory: A selective review. *Annual Review of Neuroscience*, *16*, 159–182.

Shanks, D.R., Johnstone, T., & Staggs, L. (1997). Abstraction processes in artificial grammar learning. *The Quarterly Journal of Experimental Psychology*, *50A*, 216–252.

Tunney, R.J., & Altmann, G.T.M. (1999). The transfer effect in artificial grammar learning: Reappraising the evidence of the transfer of sequential dependencies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 1322–1333.

Tunney, R.J., & Altmann, G.T.M. (2001). Two modes of transfer in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 614–639.

(RECEIVED 9/16/05; REVISION ACCEPTED 12/12/05; FINAL MATERIALS RECEIVED 1/23/06)