Journal of Experimental Psychology: General
2005, Vol. 134, No. 3, 406–419

0096-3445/05/$12.00 DOI: 10.1037/0096-3445.134.3.406

The Role of Salience in the Extraction of Algebraic Rules

Ansgar D. Endress
Ecole des Hautes Etudes en Sciences Sociales, Ecole Normale Supérieure, Centre National de la Recherche Scientifique

Brian J. Scholl
Yale University

Jacques Mehler
International School for Advanced Studies and Ecole des Hautes Etudes en Sciences Sociales, Ecole Normale Supérieure, Centre National de la Recherche Scientifique

Recent research suggests that humans and other animals have sophisticated abilities to extract both statistical dependencies and rule-based regularities from sequences. Most of this research stresses the flexibility and generality of such processes. Here the authors take up an equally important project, namely, to explore the limits of such processes. As a case study for rule-based generalizations, the authors demonstrate that only repetition-based structures with repetitions at the edges of sequences (e.g., ABCDE*FF* but not ABC*DD*EF) can be reliably generalized, although token repetitions can easily be discriminated at both sequence edges and middles. This finding suggests limits on rule-based sequence learning and new interpretations of earlier work alleging rule learning in infants. Rather than implementing a computerlike, formal process that operates over all patterns equally well, rule-based learning may be a highly constrained and piecemeal process driven by perceptual primitives—specialized type operations that are highly sensitive to perceptual factors.

Keywords: grammar acquisition, perceptual primitives, symbol manipulation, statistical learning, modularity

Cognitive processing is sometimes characterized in monolithic terms as involving a general form of symbol manipulation, association, or statistical processing. In contrast to this "one-size-fits-all" view of the mind, other researchers have characterized cognition as being an agglomeration of heuristics (e.g., Gigerenzer, Todd, & The ABC Group, 1999) or as a "bag of tricks" (e.g., Ramachandran, 1990). This debate has also been salient in research on learning in nonhuman animals, where general purpose associationist theories have been contrasted with a view in which animals are thought to have evolved specialized computational capacities for solving particular problems in their environment (e.g., Gallistel, 1990, 2000). Often the debate among these perspectives turns on how specialized and constrained the processes in question are thought to be. Here we compare such unitary versus piecemeal views of cognitive processing in the domain of rule-based generalization, in which participants must distill the abstract symbolic structure from a sequence of stimuli. Other recent research has demonstrated such abilities in both adults and young infants and has stressed the ways in which such processing may implicate a general "algebraic" symbol-manipulation architecture in the mind. In this article, in contrast, we pursue the idea that such rule-based processing may proceed using more specific primitives, for which only certain types of regularities are salient. We explore such constraints in seven experiments that collectively demonstrate how certain types of rule-based learning are constrained by *where* in the stimulus sequences the critical regularities occur. We begin by discussing some of the prominent recent research on rule-based processing in infants and its historical motivation.

Symbols and Statistics, Types and Tokens

A dominant and controversial theme in cognitive science has long been the idea that parts of cognitive processing may at root be a type of symbol manipulation in which the values of variables are stored and are subject to formal symbolic operations (e.g., Chomsky, 1980; Fodor, 1975; Fodor & Pylyshyn, 1988; Gallistel, 1990, 2000; Gallistel & Gibbon, 2000, 2002; Marcus, 2001; Newell, 1980; Newell & Simon, 1976; Pinker & Prince, 1988; Pylyshyn, 1984). This type of theorizing has met with strong resistance, however, from other researchers who view the mind as a fundamentally associative engine and who seek to explain cognitive

---

Ansgar D. Endress, Laboratoire de Sciences Cognitives et Psycholinguistique, Ecole des Hautes Etudes en Sciences Sociales, Ecole Normale Supérieure, Centre National de la Recherche Scientifique, Paris, France; Brian J. Scholl, Department of Psychology, Yale University; Jacques Mehler, Cognitive Neuroscience Sector, International School for Advanced Studies, Trieste, Italy, and Laboratoire de Sciences Cognitives et Psycholinguistique, Ecole des Hautes Etudes en Sciences Sociales, Ecole Normale Supérieure, Centre National de la Recherche Scientifique.

We are grateful to Luca Bonatti, Gennaro Chierchia, Emmanuel Dupoux, Janet Pierrehumbert, and Mohinish Shukla for helpful discussions. This research was supported by an Allocation de Moniteur Polytechnicien fellowship to Ansgar D. Endress, National Science Foundation Grant 0132444 to Brian J. Scholl, and Regione Friuli Venezia Giulia Human Frontier Science Program Grants RGP 68/2002-C and L.R. 3/98 to Jacques Mehler.

Correspondence concerning this article should be addressed to Ansgar D. Endress, who is now at the Cognitive Neuroscience Sector, International School for Advanced Studies, Via Beirut n.2-4, 34014, Trieste, Italy. E-mail: ansgar.endress@polytechnique.org

processing in terms of statistical computations such as those implemented in connectionist neural networks (e.g., Elman, 1990; Elman et al., 1996; Feldman & Ballard, 1982; Hinton & Anderson, 1981; Joanisse & Seidenberg, 1999; McClelland, Rumelhart, & The PDP Research Group, 1986; Rumelhart, McClelland, & The PDP Research Group, 1986; Seidenberg, 1997).

Recent studies of the abilities of young infants to extract various types of regularities from simple sequences have reinvigorated this debate. Saffran, Aslin, and Newport (1996), for example, showed that 8-month-old infants can use statistical information to segment a regular stream of syllables into discrete units on the basis of the frequency (and the conditional probability) with which certain syllables follow others, and they proposed that such computations may help the infant individuate words from the continuous speech signal. Others (e.g., Bates & Elman, 1996) have taken these results as an indication that statistical regularities may account for language acquisition in general.[1]

In response to this line of work, other investigators have sought to demonstrate that young infants can also extract regularities from simple sequences using formal symbolic computations, which cannot be readily accounted for in terms of association-based processing. Marcus, Vijayan, Rao, and Vishton (1999), for example, explored the abilities of 7-month-old infants to recognize structures built from various combinations of nonsense syllables such as AAB (e.g., *wo–wo–fe*), ABA (e.g., *ga–na–ga*), and ABB (e.g., *li–ti–ti*). Infants were first familiarized with a 2-min string of sequences in which each triplet was acoustically segmented from the others by a short pause. These sequences could involve several different syllables, but every triplet in the sequence corresponded to a single canonical pattern (e.g., AAB: *wo–wo–fe*, *ga–ga–na*, *li–li–ti*, etc.). After this familiarization, the infants were then tested on their ability to generalize this abstract regularity to new syllables that they had never heard during the familiarization; this was tested using the head-turn preference procedure, which measures infants' listening durations to two classes of auditory stimuli. The infants in this experiment attended longer to new test triplets that violated the previous pattern (e.g., ABA triplets, following familiarization to AAB triplets) than to test triplets that fit the previous pattern using new syllables. This pattern of results was not based simply on detecting the presence or absence of a repetition in the sequence, because the infants also discriminated AAB structures from ABB structures.

Notice that the stimuli involved in these experiments were always distinguished by the patterns of their repeated elements. In general, such repetitions can be detected in two different ways. On the one hand, participants could notice that a particular token (i.e., a particular syllable, such as *wo*) is repeated in a certain way and might later become less responsive to other presentations of this particular repetition. On the other hand, participants could instead distill the abstract structure of the repetitions from the sequence—noticing, for example, that they are hearing sequences of the type ABA, where the first syllable A is the same as the last one—independently of what the syllables A and B stand for. In this abstract case of repetition detection, A and B can be thought of as variables filled with various syllables as values. We can call this process the extraction of *type repetitions* (as opposed to merely *token repetitions*), and in this article we will always refer to this kind of type extraction when referring to *repetition-based structure*.

The infants in Marcus et al.'s (1999) experiments generalized the familiarization regularities to test stimuli involving novel syllables, and thus they clearly extracted type repetitions. These experiments are thus a powerful demonstration of early computational abilities (see Gómez & Gerken, 1999, for similar results with more complex grammars).[2] Marcus and colleagues concluded from their studies that "infants extract abstract algebra-like rules that represent relationships between placeholders (variables), such as 'the first item X is the same as the third item Y,' or more generally, that 'item I is the same as item J'" (p. 79). Under this interpretation, infants are viewed as representing serial positions as variables and are thought to be endowed with a mechanism that can discover abstract relationships among such variables.

Marcus has further argued that standard "eliminativist" connectionist networks (and other purely associative schemes) are unable to learn and generalize identity relationships in this way (e.g., Marcus, 1998a, 1998b, 2001), and these arguments have in turn been extensively criticized in the connectionist literature (e.g., Altmann, 2002; Altmann & Dienes, 1999; Christiansen, Conway, & Curtin, 2000; Christiansen & Curtin, 1999; Gasser & Colunga, 2000; McClelland & Plaut, 1999; Negishi, 1999; Seidenberg & Elman, 1999a, 1999b; Shultz & Bale, 2001; but see Marcus,

---

[1] This line of work has undoubtedly demonstrated that some "classical" nativist and symbolic theorizing underestimated the abilities of both infants and adults to learn patterns based on extremely subtle cues (e.g., Aslin, Saffran, & Newport, 1998; Newport & Aslin, 2004; Saffran, Aslin, & Newport, 1996; Saffran, Newport, & Aslin, 1996). At the same time, however, other recent work suggests that these abilities may not be fundamentally linguistic in nature: For example, they also operate over nonlinguistic auditory stimuli (Saffran, Johnson, Aslin, & Newport, 1999), over spatially variable visual stimuli (Fiser & Aslin, 2001, 2002), and in nonhuman primates (Hauser et al., 2001). Moreover, recent computational modeling suggests that segmentation using this type of statistical learning may not easily scale up to the actual speech signal involved in natural language acquisition (Yang, 2004).

[2] Gómez and Gerken (1999) found evidence for similar repetition-detection abilities in 12-month-olds using an artificial grammar learning paradigm. In short, they found that infants can generalize a finite state grammar to a new vocabulary; this generalization was probably based on specific patterns of repetitions (that did not change between vocabularies because they were characteristic of the underlying grammar). It should be noted that Gómez and Gerken (1999) did not explain their results in terms of "algebraic" rule-based processing. Instead, they proposed that the transfer could be due either to the detection of the repetition pattern, which may or may not be related to abstractions in grammar, or to "complex associations"; later experiments showed that the former interpretation is probably correct (Gómez et al., 2000; Tunney & Altmann, 2001). Gómez and Gerken (2000) contrasted the ability to generalize a repetition pattern with even more abstract category-based generalizations, in which each variable in a rule can be filled with all members of the relevant category. For example, in the (simplified) template of a transitive sentence noun–verb–noun, all nouns can fill the first and the last position; the two nouns in such a sentence need not be (and in general are not) identical. In contrast to the ABA pattern in Marcus et al.'s (1999) experiments where individual syllables had to be repeated, what is repeated in the noun–verb–noun template are the *categories* to which the items in the sentence belong. Such category-based generalizations appear to be required in natural-language processing and are not fully captured by studies in which variables must be filled with particular physical tokens (such as the syllables in Marcus et al.'s, 1999, experiments).

1999a, 1999b, 1999c, 1999d, 1999e). We will not enter into this debate here, though we believe that Marcus (1999a, 1999b, 1999c, 1999d, 1999e) has put forward principled and compelling arguments in this ongoing controversy. Rather, in this article, we focus on the various types of processing that could underlie the extraction of symbolic structure such as that involved in Marcus et al.'s (1999) experiments. In particular, we ask how general the responsible processing must be: Do such processes operate over all regularities equally in the manner of formal operations in a computer, or might such processing be constrained by the relative perceptual salience of various parts of the familiarization sequences?

Are Repetition-Based Structures Diagnostic of a General Rule-Extraction Process?

Although the infants in the experiments reported above clearly generalized the repetition-based structures, there is still an important ambiguity as to what is meant by the term *algebraic rule.* On the one hand, this term can refer to an abstract formal relation among variables, which is fully general in the sense that it can operate over any regularity equally well. (Such general abilities may be available to human adults, because the fixation of belief is thought to be analogous to scientific inference; e.g., Fodor, 1983). One can easily imagine a simple operation of this type in a digital computer, for example, which detects the presence of a repetition equally well in a sequence of a given length, regardless of just how and where the repetition occurs inside the sequence. Indeed, many of the underlying operations implemented in computer chips do in fact work this way—for example, shifting a register by one place works equally efficiently regardless of which register has to be shifted and regardless of its content. In fact, many authors have proposed that the powerful human problem-solving capacities result from general symbolic operations that may be modeled with symbol manipulation in a computer (e.g., Anderson, 1989, 1993; Lehman, Laird, & Rosenbloom, 1998; Newell, 1980; Newell & Simon, 1976).

On the other hand, an algebraic rule could be thought of in more piecemeal terms, as a Gestalt-like operator that could efficiently detect repetitions of only particular sorts, based on the perceptual salience of the repetition locations. Note that both interpretations here are abstract—that is, they both represent the repetition structure independently of the particular physical stimuli by which it is carried. Nevertheless, this difference is a critical one for assessing how demonstrations of symbolic rule-learning may scale up: The first view suggests that such abilities may hold for the extraction and generalization of any computable regularity, whereas the second view suggests that we may find successful generalization only in highly specific situations.

Which of these interpretations is implicated by the demonstrations of infants' symbolic rule-learning discussed above? Surprisingly, the authors of these studies do not address this issue directly, even in later attempts to explore the wider theoretical implications of these results (Marcus, 2001; Marcus et al., 1999). In general, however, these studies have been interpreted in ways that are entirely consistent with the fully general and formal view. Marcus (2001), for example, suggested that registers and operations in computers might be good models of variables and operations in the human mind, and he discussed how registers could be implemented in neural tissue. Although he suggested that the mind contains only a limited stock of basic algebraic operations, his discussion suggested that the operations which do exist operate equally well over all types of input in the manner of such operations in digital computers:

> Variables are one part of the story, operations over those variables are another. To clarify the difference, consider the distinction in digital computers between registers and instructions. Registers store values; instructions, such as "copy" and "compare" manipulate those values. My hunch is that the brain contains a similar stock of basic instructions, each defined to operate *over all possible values of registers* [italics added]. (Marcus, 2001, p. 58)

In emphasizing the general applicability of such operations, Marcus cleaved to how such operations actually work in digital computers—wherein a register's contents are copied or compared in exactly the same way (and with the same reliability) regardless of their contents. Beyond a few passages of this sort, however, previous work has simply remained silent about how general such symbolic operations may be, and our goal in the present article is to address this issue empirically.

Although this contrast represents an open empirical question, there are several reasons to suspect that there might be lower level primitive mechanisms for repetition detection in the human mind. For example, it has been shown that the ability to generalize visually presented sequences constructed using finite state grammars (artificial grammars commonly used to study implicit learning) depends on characteristic repetition patterns (Gómez, Gerken, & Schvaneveldt, 2000; Tunney & Altmann, 2001).

In the present experiments, we focus on the position of repeated elements in the spontaneous generalization of auditory sequences. The structures used in the experiments of Marcus et al. (1999)—that is, AAB, ABA, and ABB—not only involved repetitions but all repetitions were always located at either the initial or final edge of the sequence.[3] Edge positions of this type are known to be salient in many ways. In most natural languages, for example, poetic rhyme is defined with respect to the last syllable of a word, and there are no analogous stylistic figures defined with respect to the middle of a word. More generally, it has been known since Ebbinghaus (1885/1913) that edge positions are more prone to be remembered than internal positions in a sequence, a finding that has been termed the *serial position effect*. The repetition-based structures in Marcus et al.'s experiment were therefore perceptually highlighted, and this might have enabled their detection and generalization even by a highly constrained "perceptual primitive" operation, which was sensitive only to perceptually salient repetition structures.

---

[3] The term *edge* simply refers to a sequence-initial or sequence-final structure. We use this term here to capture the sense that there is something salient and perceptually nonarbitrary about such positions. At the same time, however, we do not mean to entail any other similarities with spatial edges as they are studied in early visual processing.

The Current Experiments: Adults' Spontaneous Generalization of Repetitions of Differing Perceptual Salience

In the experiments reported below, we explore the general applicability of "repetition detection" in symbolic generalization by contrasting repetitions that occur at the final edges of sequences (i.e., ABCDE*FF*) and those that occur internally (i.e., ABC*DD*EF). A fully general rule-extraction mechanism—as in a digital computer—should operate equally efficiently over both sorts of repetitions. In contrast, a more piecemeal process specialized for the detection of repetitions might only operate over salient positions such as edges, suggesting that participants may not be able to efficiently generalize sequence-internal repetitions.

Note that in order to carry out these experiments, we had to resort to longer sequences than those used by Marcus et al. (1999)—simply because the use of triplets necessarily entails that repetitions will occur at an edge if they occur at all. (The only two repetition-based structures that are available in triplets are AAB and ABB.) In this sense, the stimuli used in these previous studies did not even allow for the possibility of exploring the primary question addressed in our experiments. However, because these previous studies tested infants, they cannot be faulted for cleaving to the shortest strings possible. In general, experiments on such populations have revealed that the indirect measures used to test the underlying competence are easily overwhelmed by performance constraints and that making stimuli more complex will often distract the participants from the regularities of interest. Thus, it seems likely that infants, for example, might fail to generalize symbolic structures from longer sequences simply because of such performance limitations. In our experiments, we thus tested spontaneous symbolic generalization in adults. This would be nonoptimal in a situation aimed at demonstrating an underlying competence, because infants cannot be expected to have all of the mature abilities of adults, in any domain. Here, however, we are in search of *constraints* on such processing, and there is no reason to think that the generalization abilities of young infants would be any less prone to perceptual constraints than those of adults.

In Experiment 1a, we tested participants' ability to spontaneously generalize the structure of auditory syllable sequences of the form ABCDE*FF*—where the critical generalization always involved the final repeated syllable of each sequence. Experiment 1b was identical to Experiment 1a except that participants had to spontaneously generalize the structure of sequences of the form ABC*DD*EF—where the critical repetitions were now sequence-internal. In both of these experiments, participants were given a cover task to ensure that they attended to the sequences, but this cover task did not prioritize any particular elements. In Experiments 2a and 2b, we then attempted to replicate Experiments 1a and 1b under mere exposure conditions in which participants had no task at all during familiarization (except to listen to the stimuli). Finally, in Experiments 3a–3c, we asked whether the advantage for edge-based repetitions in spontaneous generalization could be due to a general failure to process the sequence-internal syllables. In Experiment 3a, participants had to discriminate sequences of the form ABCDE*FF* from sequences of the form ABCD*EE*F. In Experiments 3b and 3c, participants had to discriminate sequences of the form ABC*DD*EF from sequences of the form AB*CC*DEF. Successful discrimination in both conditions would suggest that any advantage for edge repetitions in the earlier experiments is specific to rule-based generalization, rather than reflecting a more general perceptual advantage.

Experiment 1: Generalizing Sequence-Final and Sequence-Internal Repetitions

In Experiment 1, we asked whether participants' ability to generalize repetition-based structures would depend on the location of the repetitions. In Experiment 1a, the repetition was located at the final edge of sequences, whereas it was sequence internal in Experiment 1b.

Participants

Forty participants (30 women, 10 men, mean age = 23.3 years, range = 19–37 years) were randomly assigned to Experiments 1a and 1b as described below. All participants were native speakers of Italian. In these and all other experiments, participants reported normal or corrected-to-normal vision, normal audition, and no history of neurologic and psychiatric disorders.

Experiment 1a: Generalizing Sequence-Final Repetitions

We first tested participants' ability to spontaneously generalize repetition-based structures when the repetition was located at the final edge of each sequence. We familiarized participants with 36 exemplars of syllable sequences of the form ABCDE*FF* (e.g., /zOfesapitu*koko*/, with O being the vowel in the French word *deux*, meaning two); subsequently, participants had to decide whether novel sequences of the form ABCDE*FF* or ABCD*EE*F (e.g., /zOfesapi*tutu*ko/) contained the same regularity as the sequences with which they had been familiarized.

Method

Apparatus. Participants were tested individually in a quiet room; the experiment was run on a PC using Presentation (Neurobehavioral Systems, n.d.) software, and all stimuli were presented over headphones.

Stimuli. We synthesized the stimuli with the *fr2* (French female) diphone base of MBROLA (Dutoit, Pagel, Pierret, Bataille, & van der Vreken, 1996). We used exclusively consonant–vowel syllables with phoneme durations of 116 ms and a first formant frequency of 200 Hz;[4] the syllable duration was therefore 232 ms. There was no silence among syllables within a sequence.[5] We used the consonants *p*, *t*, *k*, *z*, *f*, and *s* and the vowels *a*, *e*, *i*, *o*, *u*, and O. We used different sequences in the familiarization phase and in the test phase; participants could therefore not

---

[4] More precisely, we used a pitch point of 200 Hz at 50% of the segment duration.

[5] Peña et al. (2002) showed that imperceptible silences in an otherwise continuous speech stream can induce generalizations that are not otherwise available. Because this kind of generalization could have biased our results, we used continuous syllable sequences without silences.

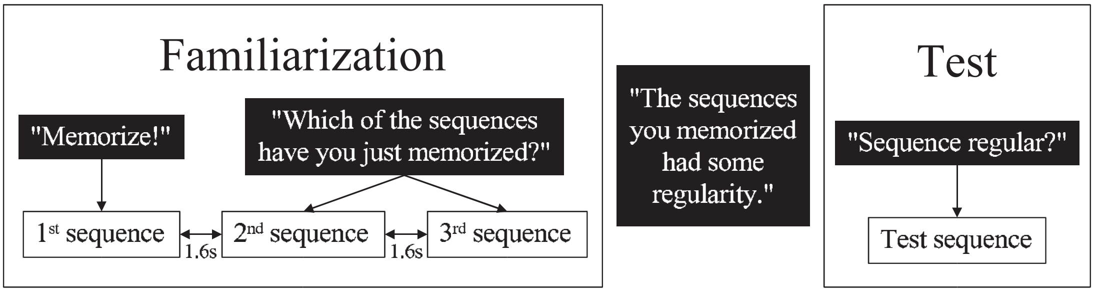

Figure 1. Trial structure in Experiments 1a and 1b. White-on-black rectangles represent instructions that appeared on the screen, whereas black-on-white rectangles represent the stimuli.

succeed by using transition probabilities as in statistical learning experiments.[6]

Familiarization. We familiarized participants with syllable sequences of the structure ABCDE*FF*. Participants were instructed to attend to the familiarization sequences and to commit them to memory as well as they could. To reduce variability among participants and to evaluate whether participants paid attention to the sequences, we tested whether participants retained each sequence immediately after they heard it. More specifically, participants were presented with three sequences in a row. They were instructed to memorize the first sequence and then to decide whether it was identical to the second or the third sequence (see Figure 1).[7] The first sequence always had the structure ABCDE*FF*. One of the two following sequences was identical to the first one; the other contained the same syllables in the same order, but its fourth syllable was repeated (i.e., ABC*DD*EF). Sequences within a trial were separated by a silence of 1,624 ms (i.e., one sequence duration). Trials were separated by a 2-s delay after each response. The familiarization phase included 36 trials.

Test. After the familiarization phase, participants were informed that the sequences they had just encountered had contained some regularity. They were then told that they would now hear new sequences and would have to decide whether these sequences conformed to the same regularity as the sequences they had been asked to memorize during familiarization.[8] They were presented with 18 new sequences that conformed to the structure ABCDE*FF* (as did the familiarization trials), mixed with another 18 new sequences that conformed to the structure ABCD*EE*F (an entirely novel structure). All test sequences used an entirely new set of syllables, which were not used during familiarization. We will call a sequence *grammatical* if it conformed to the structure participants had been familiarized with and *ungrammatical* otherwise. Trials were separated by a 2-s delay after each response. The structure of the trials during familiarization and during test is summarized in Figure 1. Participants were never provided with feedback in any of the experiments reported here.

Results

The mean percentage of correct responses during familiarization was 86.6% ± 10.9% (mean percentage of correct responses ± standard deviation), which was reliably better than chance, $t(22) = 16.1$, $p < .001$. Participants who made more than 20% errors during familiarization were excluded from further analyses, yielding 15 participants in the test phase. (Although this threshold was arbitrary, the motivation for the retention test during familiarization was precisely to detect participants who did not attend to the stimuli. Nevertheless, as we will show below, our results are not qualitatively altered when including all participants—or, for that matter, when choosing any other threshold.) As depicted in Figure 2, participants correctly generalized the structure of the memorized sequences to new sequences (70.7% ± 20.2%; Cohen's $d = 1.06$), $t(14) = 4.0$, $p = .001$.[9]

Experiment 1b: Generalizing Sequence-Internal Repetitions

The participants in Experiment 1a successfully generalized the structure ABCDE*FF*, in which the critical repetition occurred at the final edge of each sequence. In this experiment, we asked whether participants would also succeed in generalizing repetition-based structures when the repetition was located in the middle of a sequence, as in ABC*DD*EF.

---

[6] A list of all used stimuli is available from http://www.ehess.fr/centres/lscp/persons/endress/publications/Salience_and_Rule-Extraction-Appendix.pdf.

[7] Participants were given the following instructions (translated from Italian):

> Hello, and thanks for choosing to participate! This experiment consists of two parts. In the first part, you are going to hear 36 groups of 3 seven-syllable sequences. Try to memorize the first sequence. The second or the third sequence will be identical to the first one. After having listened to the three sequences, you should decide whether the second or the third sequence was identical to the first one. Press 2 if you think that the second sequence was the same as the first one. Press 3 if you think that the third sequence was identical to the first one. If you are not sure, trust your intuition. Press 1 to continue.

[8] Participants were given the following instructions (translated from Italian):

> Welcome to the second part of the experiment! The sequences that you have memorized in the first part contained a regularity. Now you are going to hear new seven-syllable sequences. Half of the sequences contain the same regularity as the sequences you have memorized in the first part. Press 1 if you think that a sequence contains the same regularity as the sequences you have already memorized. Press 2 if you think that a sequence does *not* contain the same regularity as the sequences you have already memorized. If you are ready, press 1 to continue.

[9] When all participants were included, they still succeeded in generalizing the structure (61.8% ± 22.4%; Cohen's $d = 0.54$), $t(22) = 2.5$, $p = .02$.

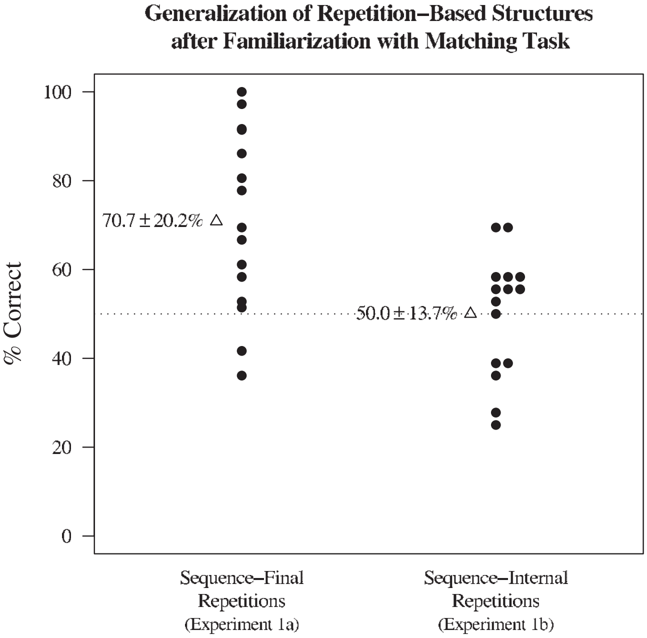

Figure 2. Results of the generalization tests in Experiments 1a and 1b. Participants generalized the structure ABCDE*FF* but not the structure ABC*DD*EF. Values represent the mean percentage of correct responses ± its standard deviation. Each dot represents a participant. The dotted line represents the chance level of 50%.

Method

This experiment was identical to Experiment 1a except that new types of sequences were used. Participants heard sequences of the structure ABC*DD*EF during the familiarization phase. Each sequence was followed by an identical sequence and by a sequence that contained the same syllables in the same order but conforming to the structure ABCDE*FF*. Participants then had to choose which of the two consecutive sequences was the one they had previously heard. During the test phase, participants were presented with new sequences conforming to the structures ABC*DD*EF (grammatical) and AB*CC*DEF (ungrammatical).

Results

The mean percentage of correct responses during familiarization was 92.0% ± 8.4%, $t(16) = 20.7$, $p < .001$; this score was not different from that observed in Experiment 1a, $F(1, 38) = 2.9$, $p = .096$, *ns*. For the test phase, we excluded participants from the analysis who made more than 20% errors during familiarization. Fifteen participants were thus included in the analysis of the test phase.

As is clear from inspection of Figure 2, participants failed to generalize the structure ABC*DD*EF to new sequences in the test phase; their percentage of correct responses did not differ from chance (50.0% ± 13.7%; Cohen’s $d = 0.00$), $t(14) = 0.0$, $p > .999$, *ns*.[10] They failed both to endorse grammatical sequences (51.5% ± 23.6%), $t(14) = 0.2$, $p > .80$, *ns*, and to reject ungrammatical sequences (48.5% ± 23.1%), $t(14) = 0.3$, $p > .80$, *ns*. The difference between grammatical and ungrammatical sequences was due to 1 participant who endorsed virtually all sequences as grammatical and failed to reach significance (Cohen’s $d = 0.08$), $t(14) = 0.3$, $p > .70$, *ns* (paired *t* test). There was no difference between the percentages of endorsements as grammatical in the first 18 trials (51.1% ± 20.3%) and the last 18 trials (51.9% ± 18.3%; Cohen’s $d = 0.08$), $t(14) = 0.4$, $p > .70$, *ns* (paired *t* test). The percentage of correct responses in the test phase of Experiment 1b was significantly different from that in Experiment 1a (Cohen’s $d = 1.33$), $F(1, 28) = 10.7$, $p = .003$.[11]

Discussion

Participants succeeded at generalizing repetition-based structures when the repetitions were located at the final edge of the sequences (Experiment 1a), but they failed to generalize such structures when the repetition was sequence-internal (Experiment 1b). Below we interpret this contrast in terms of the operation of a *perceptual primitive* for repetition detection that is constrained by the perceptual salience of the various repetition positions within a sequence. First, though, we consider some other possible interpretations of these results.

Perhaps participants failed to generalize sequence-internal repetitions simply because they had learned that a repetition was *somewhere* in the middle of a sequence without knowing exactly where. In other words, participants might well have noticed the repetitions but may have had a greater difficulty in tracking their exact position in the sequence. It is important to note that in principle, this result could fit comfortably with the piecemeal approach sketched in the introduction: This is simply one way for participants to fail, in a situation where this possibility was no more formally likely—and no more predicted by a fully general rule-extraction mechanism—than would be a failure in Experiment 1a for the same reason. Indeed, from a formal point of view, participants may just as well have learned only that a repetition was “somewhere near the edge of a sequence, without knowing exactly where.” In fact, however, aspects of our data speak against this interpretation. If participants had detected the sequence-internal repetitions but had difficulty in tracking their exact positions, then they should have had a bias toward considering all sequences as grammatical, because all sequences had a repetition somewhere in their middle. This possibility was disconfirmed, however: Participants failed to accept grammatical sequences just as often as they failed to reject ungrammatical sequences.

A more nuanced version of this interpretation might take into account the additional possibility that participants rejected some sequences simply because they knew that the test phase would contain both grammatical and ungrammatical sequences. This view, however, still predicts that participants should start out by endorsing all sequences as grammatical (because they all contained a sequence-internal repetition) and should switch to rejecting some sequences only when they notice that they have not yet seen any ungrammatical sequences. This prediction was also not

---

[10] When all participants were included, participants still failed to generalize the structure ABC*DD*EF (49.3% ± 13.0%; Cohen’s $d = -0.05$), $t(16) = 0.2$, $p > .80$, *ns*.

[11] The percentage of correct responses was also different between the two experiments when all participants were included (Cohen’s $d = 0.68$), $F(1, 38) = 4.2$, $p = .048$. Excluding participants who failed to remember the familiarization sequences simply had the effect of reducing variance between participants due to attentional factors.

confirmed, because participants were no more likely to endorse sequences as grammatical in the first half of the test phase compared with the last half.

Finally, we question whether the familiarization task may have simply biased participants to process edges more than sequence-internal syllables: They could have succeeded on the familiarization task, in principle, simply by detecting whether a repetition was located at the final edge versus anywhere else. During the familiarization of Experiment 1b, participants had to memorize sequences conforming to structure ABC*DD*EF. Each sequence was followed by an identical sequence and by a sequence that contained the same syllables in the same order but conforming to the structure ABCDE*FF*. Then, participants had to choose which of the two consecutive sequences was the one they had previously memorized. Participants could have noticed that one of these two sequences had a repetition in its final edge and that this sequence was unlike the sequence they had memorized (which had no repetition in its final edge). This situation may have led to more attention paid to the edges of sequences during familiarization and thus to success at generalizing edge repetitions but failure in generalizing sequence-internal repetitions.

The critical thing to say about this possibility is that its underlying logic is intuitive but misleading. In fact, there was nothing about the familiarization task that emphasized edges over internal positions: By this logic, participants could just as easily have succeeded in the familiarization task of Experiment 1a simply by detecting whether a repetition was located "in the fourth position versus anywhere else." Of course, the intuition here is that edges are intrinsically more salient than particular internal positions—that is, that participants are more likely to confuse one internal position with another than they are to confuse an edge position (e.g., ABCDE*FF*) with a near-edge position (e.g., ABCD*EE*F). In fact, however, this is exactly the sort of phenomenon that we are exploring in these experiments, because this difference does not in any way fall out of the formal structure of the sequences. In a typical operation of a digital computer, for example, neither of these distinctions would be more salient than the other, simply because the salience of particular sequence positions does not play a role in the first place: Operations such as "shift," "compare," or "copy" would rather treat all sequences equally. Nevertheless, we also ruled out this possibility experimentally in Experiments 2a and 2b.

Experiment 2: Generalizing Sequence-Final and Sequence-Internal Repetitions With Simple Exposure

Above we considered the possibility that the familiarization tasks used in Experiments 1a and 1b may have inadvertently biased participants toward the edges of the sequences, but we rejected this possibility on theoretical grounds. We now also seek to rule out this hypothesis experimentally in the most direct way possible (while simultaneously replicating our primary result) simply by repeating these experiments with no familiarization task at all. Participants simply listened to the sequences during the familiarization phase, with no task except to memorize them to the best of their ability.

Participants

Thirty participants (15 women, 15 men, mean age = 24.6 years, range = 20–37 years) were randomly assigned to Experiments 2a and 2b. All were native speakers of French.

Experiment 2a: Generalizing Edge Repetitions With Simple Exposure

Method

This experiment was identical to Experiment 1a except that participants had no task during familiarization except to attend to the sequences and to commit them to memory as well as they could. The sequences in the familiarization phase all conformed to the structure ABCDE*FF*, and participants proceeded through the sequences at their own pace by pressing a key after each one.

Results

As depicted in Figure 3, participants correctly generalized the structure from the memorized sequences to new sequences during the test phase (82.8% ± 19.8%; Cohen's $d = 1.71$), $t(14) = 6.4$, $p < .001$. This experiment thus replicated the primary pattern of results from Experiment 1a.

Experiment 2b: Generalizing Sequence-Internal Repetitions With Simple Exposure

Method

This experiment was identical to Experiment 1b except that, as in Experiment 2a, participants had no task during familiarization except to

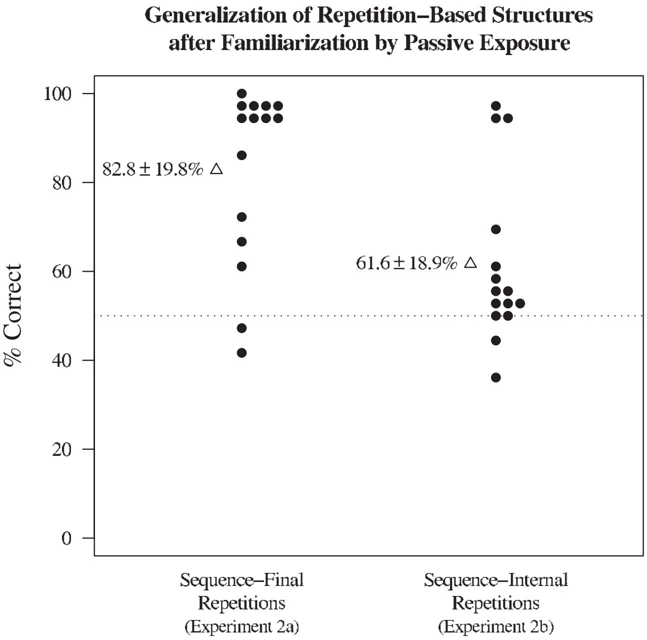

Figure 3. Results of Experiments 2a and 2b. Participants generalized the structure ABCDE*FF*, but the overwhelming majority failed to generalize the structure ABC*DD*EF. Values represent the mean percentage of correct responses ± its standard deviation. Each dot represents a participant. The dotted line represents the chance level of 50%.

attend to the sequences and to commit them to memory as well as they could. The sequences in the familiarization phase all conformed to the structure ABC*DD*EF, and participants proceeded through the sequences at their own pace by pressing a key after each one.

Results

Participants generalized the structure to new sequences (61.6% ± 18.9%; Cohen's $d = 0.64$), $t(14) = 2.4$, $p = .032$. They endorsed grammatical sequences as grammatical (70.0% ± 26.4%), $t(14) = 2.9$, $p = .011$, but failed to reject ungrammatical sequences as ungrammatical (53.3% ± 31.1%), $t(14) = 0.4$, $p > .60$, *ns*. However, this difference failed to reach significance (Cohen's $d = 0.40$), $t(14) = 1.4$, $p = .160$, *ns* (paired *t* test). Participants did not exhibit an overall bias toward endorsing or rejecting sequences (58.3% ± 21.8% endorsements as grammatical), $t(14) = 1.4$, $p = .160$, *ns*. There was no difference between the percentages of endorsements as grammatical in the first 18 trials (58.1% ± 23.0%) and the last 18 trials (58.5% ± 23.3%; Cohen's $d = -0.02$), $t(14) = 0.1$, $p > .90$, *ns* (paired *t* test).

As is clear from inspection of Figure 3, however, this overall pattern of results is misleading. In particular, the relatively high mean performance in this experiment was due to 3 participants who discovered the structure explicitly and made at most two mistakes. (This was probably possible because participants were given the opportunity to rehearse each training sequence for as long as they wanted before advancing to the next sequence.) Even so, participants performed considerably better in Experiment 2a than in Experiment 2b (Cohen's $d = 1.13$), $F(1, 28) = 8.9$, $p = .006$. Excluding the 3 participants who had extracted the structure explicitly, the percentage of correct responses was only 53.2% ± 8.3% (Cohen's $d = 0.41$), $t(11) = 1.6$, $p > .20$, *ns*, with most participants replicating the failure observed in Experiment 1b.[12]

Discussion

Experiments 2a and 2b collectively replicate the main trends from Experiments 1a and 1b: Most participants readily and spontaneously generalized structures containing edge repetitions (i.e., ABCDE*FF*) but failed to generalize structures containing sequence-internal repetitions (i.e., ABC*DD*EF). The overall failure with sequence-internal repetitions thus rules out the possibility that the familiarization task used in Experiments 1a and 1b could explain the overall pattern of results—because neither Experiment 2a nor 2b had any familiarization task at all except to listen to the familiarization sequences. Indeed, the contrast between performance in Experiments 2a and 2b is particularly instructive given that participants had the possibility to rehearse each sequence ad libitum. Even under these conditions, only 3 participants distilled the abstract structure ABC*DD*EF, and all of these 3 participants had not only understood the structure of the sequences but were also able to report the algorithm with which the sequences had been generated. When asked whether they had noticed the structure of the sequences (a standard debriefing question), the participants reported the sequence-internal repetitions as well as the details of how vowels and consonants had been assigned to the syllables in the training and test sequences. We conclude that these 3 participants represented outliers not only in the quantitative sense (far outperforming all other participants in this experiment) but also in the conceptual sense that they solved the task explicitly, using a different kind of rule-extraction process than the other participants. (We want to be careful here not to make any particular claims as to what can be learned explicitly or implicitly in general. We simply note that our informal debriefings never generated any other suggestion—except for these 3 outliers—that any participant in Experiments 1b or 2b had explicitly learned the nature of the manipulation.) Overall, the pattern of results in this experiment suggests that people are capable of entertaining and testing complex hypotheses when given the opportunity to verbalize them but that unless they notice the relevant structure explicitly, all sequences are not treated equally: Edge-based repetitions will be spontaneously generalized more readily than will sequence-internal repetitions.

Experiment 3: Generalization Versus Discrimination

In interpreting the previous experiments, we have suggested that the process underlying symbolic generalization may prioritize sequence edges over sequence-internal positions. Another (and more deflationary) possible explanation for this pattern of results, however, is that participants may simply have a more fundamental difficulty with accurately encoding items in sequence-internal positions. This view would predict, however, that sequence edges should be prioritized not only when participants have to generalize the abstract structure of sequences beyond particular tokens but also when they must simply *discriminate* sequences. In our final three experiments, therefore, we constructed a task in which participants had to make the same discriminations as in the test phases of Experiments 1a and 1b but in which generalization to new tokens in the test phase was not required. In this first discrimination experiment, participants simply made same–different judgments on pairs of the same types of test sequences used in the previous generalization experiments, discriminating edge repetitions (ABCDE*FF* vs. ABCD*EE*F).

Participants

Thirty participants (21 women, 9 men, mean age = 25.5 years, range = 20–63 years) were randomly assigned to Experiments 2a and 2b. All were native speakers of French or Italian.

Experiment 3a: Discriminating Edge Repetitions

Method

This experiment was identical to the previous experiments except as noted here. In each trial, participants heard two syllable sequences that could be either the same or different. The first sequence always conformed to the structure ABCDE*FF*. The second sequence was either identical to the first one or was a sequence with the same syllables in the same order but conforming to the structure ABCD*EE*F. Participants were instructed to

---

[12] As is clear from inspection of Figure 3, the results of Experiment 2a did not contain any outliers of the type that influenced this experiment. Even when the 3 best performing participants are removed from Experiment 2a, successful rule generalization is still highly statistically robust (79.0% ± 20.4%), $t(11) = 4.9$, $p < .001$: The remaining participants from Experiment 2a performed much better than participants in Experiment 2b, including the 3 outliers (Cohen's $d = 0.88$), $F(1, 25) = 5.2$, $p = .032$.

respond with a button press on each trial to indicate whether the two sequences were the same or different. They thus had to make the same distinctions as in the test phases of Experiments 1a and 2a but now without the need to draw abstract generalizations. The sequences were randomly chosen from those used in the test phase of Experiment 1a.

To make these studies as similar as possible to the previous generalization experiments, we also included a familiarization phase. Before proceeding to the actual trials, participants were first presented with 36 example trials (18 in which the correct response was "different" and 18 in which the correct response was "same"). The sequences for these examples were randomly chosen from those that were used in the familiarization phase of Experiment 1a.[13]

Results

As is clear from Figure 4, discrimination performance was near perfect (93.7% ± 5.4%; Cohen's $d = 8.38$), $t(14) = 31.3$, $p < .001$.

Experiment 3b: Discriminating Sequence-Internal Repetitions

Method

This experiment was identical to Experiment 3a except for three changes: (a) The structure ABCDE*FF* was replaced by ABC*DD*EF, (b) the structure ABCD*EE*F was replaced by AB*CC*DEF (i.e., with a change to the repetition location in the same direction as was used in the previous experiment), and (c) all sequences were chosen from the sequences of Experiment 1b. In other words, participants were presented with pairs of sequences, the first of which always conformed to the structure ABC*DD*EF. The second sequence was then either identical to the first one or was a sequence with the same syllables in the same order but conforming to the structure AB*CC*DEF. Participants had to decide whether the sequences were the same or different. Before proceeding to the actual test trials, participants were presented with 36 example sequences as in Experiment 3a. If the difference in generalization performance between the structures ABCDE*FF* and ABC*DD*EF in the previous experiments was due to a more fundamental deficit for processing items located in the interior of a sequence, then participants should perform worse in this experiment than in Experiment 3a.

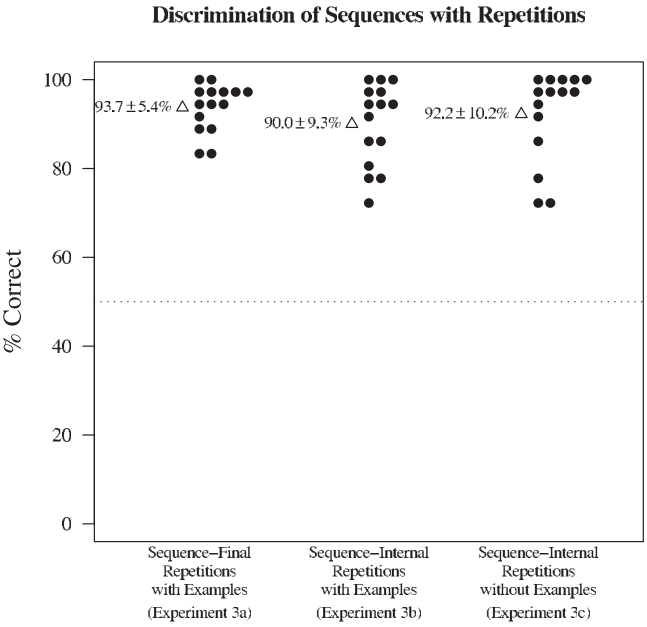

Figure 4. Results of Experiments 3a–3c. Participants discriminated sequences conforming to the structure ABCDE*FF* from sequences conforming to the structure ABCD*EE*F as well as sequences conforming to the structure ABC*DD*EF from sequences conforming to the structure AB*CC*DEF. In Experiments 3a and 3b (left and middle, respectively), participants were first presented with example sequences but not in Experiment 3c (right). Values represent the mean percentage of correct responses ± its standard deviation. Each dot represents a participant. The dotted line represents the chance level of 50%.

Results

As depicted in Figure 4, the discrimination performance was near perfect (90.0% ± 9.3%; Cohen's $d = 4.47$), $t(14) = 16.7$, $p < .001$. Performance was not different from that obtained in Experiment 3a (Cohen's $d = 0.56$), $F(1, 28) = 1.8$, $p = .192$, *ns*. Participants thus discriminated the sequences in Experiments 3a and 3b equally well, suggesting that the previous failure to generalize the abstract structure of ABC*DD*EF was not due to lower level psychophysical difficulties.

Experiment 3c: Discriminating Sequence-Internal Repetitions Without Previous Exposure

In Experiments 3a and 3b, we tested participants' ability to make simple same–different discriminations using sequences that had either edge repetitions or sequence-internal repetitions. In these studies, we included a familiarization phase in an attempt to make these experiments as similar as possible to the earlier studies of abstract generalization (Experiments 1a–2b). However, it is possible that the example sequences used in Experiment 3b might have somehow encouraged participants to attend more carefully to the sequence-internal positions than did the familiarization sequences from the generalization experiments (either the familiarization tests used in Experiments 1a and 1b or the simple exposure familiarization used in Experiments 2a and 2b). To rule out this possibility, we attempted to replicate the successful discrimination for sequence-internal repetitions in this experiment, without any example sequences at all.

Method

Experiment 3c was identical to Experiment 3b except that the example sequences were not included; participants simply proceeded directly to the test phase.

Fifteen participants were tested (9 women, 6 men, mean age = 23.4 years, range = 18–35 years). All were native speakers of French.

---

[13] Participants were given the following familiarization instructions (in French or Italian):

> Hello, and thanks for having decided to participate. You are going to hear 36 pairs of seven-syllable sequences. You will have to decide whether the two sequences are the same or different. Before starting the experiment, you are going to here 36 *examples*. Try to compare the sequences in each pair in order to see the difference if there is one. Press 1 to continue.

Results

As depicted in Figure 4, the discrimination performance was again near perfect (92.2% ± 10.2%; Cohen's $d = 4.29$), $t(14) = 16.0$, $p < .001$, and was not different from the performance obtained in either Experiment 3a (Cohen's $d = 0.22$), $F(1, 28) = 0.2$, $p = .623$, *ns*, or Experiment 3b (Cohen's $d = 0.24$), $F(1, 28) = 0.4$, $p = .537$, *ns*. The proportion of participants who gave a correct response on the first trial was different from chance level ($p < .008$, binomial test), and the percentage of correct responses on each participant's first trial with identical sequences was not different from the percentage on each participant's first trial with different sequences, $\chi^2(1, N = 30) = 2.16$, $p > .141$, *ns*.

Discussion

In Experiments 1a–2b, participants readily generalized the structure ABCDE*FF* but failed to generalize the structure ABC*DD*EF. In contrast, Experiments 3a–3b showed that participants are able to discriminate sequences with repetitions equally well, regardless of whether the repetitions occur at an edge or sequence-internally.

Further analyses showed that participants performed well on their first trial and that the initial performance for trials with identical and different sequences did not differ. These findings suggest that participants were not encouraged by the experiment to attend specifically to sequence-medial positions but rather they could spontaneously encode and process syllables in these positions.[14]

There are other ways in which the discrimination task in Experiments 3a–3c might have been simpler than the generalization task in Experiments 1a–2b. For example, in Experiments 3a–3c, participants had to remember only two sequences at a time, whereas they had to memorize all familiarization sequences before being tested in the other experiments. Of course, they could not possibly have memorized all the sequences in the generalization experiments because they were simply too many of them. More important, even if they had memorized them, it would not have facilitated the generalizations, because participants were tested on new sequences. The difference between Experiments 3a–3c and Experiments 1a–2b thus cannot be attributed only to memory demands, suggesting that the failure to generalize the structure ABC*DD*EF was not due to a brute processing impairment for sequence-internal tokens; rather, the difficulty was specific to generalizations.

General Discussion

The idea that various aspects of cognition involve formal symbolic operations over variables is arguably one of the most important and controversial themes in cognitive science (e.g., Chomsky, 1980; Fodor, 1975; Fodor & Pylyshyn, 1988; Gallistel, 1990, 2000; Gallistel & Gibbon, 2000, 2002; Marcus, 2001; Newell, 1980; Newell & Simon, 1976; Pinker & Prince, 1988; Pylyshyn, 1984). Such theorizing has recently been invigorated by experimental evidence demonstrating that both infants and adults can spontaneously generalize symbolic rules from simple sequences (e.g., Gómez & Gerken, 1999; Marcus et al., 1999; Peña, Bonatti, Nespor, & Mehler, 2002). The experiments reported here explored the nature of such symbolic generalizations and contrasted two possibilities that had not been distinguished in previous work.

On the one hand, such generalization could reflect a completely general formal process, operating over any regularity equally well. In other words, such operations could be realized in the mind in the same manner in which they are realized in a digital computer; that is, they could consist in, for example, shifting, copying, or comparing registers regardless of their contents. On the other hand, rule-based generalizations could reflect more constrained and piecemeal processes; they might be mediated by specific operators that are able to detect only particular structures, based on their perceptual salience or the degree to which they match Gestalt-like templates.

Generalizations in Edges and Middles

The experiments reported here used repetition-based structures as a case study to contrast these two possibilities, and the results supported the piecemeal interpretation. We examined the ability of human adults to process repetition-based structures whose repetitions were located in various positions within simple auditory syllable sequences. Participants were familiarized with sequences that had a repetition in a characteristic position and had to extract this abstract structure, as defined by the sequential position where the repetition occurred. We asked whether participants could generalize such structures to new sequences with new syllables. In Experiments 1a–2b, participants generalized the structures only when the repetitions were located at the edges of sequences: They readily generalized the structure ABCDE*FF* but failed to generalize the structure ABC*DD*EF. This pattern of results was observed when a familiarization task ensured that participants paid attention to the familiarization sequences (in Experiments 1a and 1b) and also under "mere exposure" conditions wherein participants simply listened to the familiarization sequences without any other task (in Experiments 2a and 2b). In no case did participants have any knowledge of just how they would be later tested.

Critically, the contrast between edges and sequence-internal positions occurred only for symbolic generalization and did not reflect a universal inability to process the sequence-internal positions: Experiments 3a–3c demonstrated that participants were able to discriminate sequences with repetitions equally well, regardless of whether the repetitions occurred at an edge or in a sequence-internal position. This argues against the possibility that our results reflect a generic performance difficulty for processing items in sequence-internal positions, which affects the encoding of the corresponding syllables. That such constraints do not affect the processing of sequence-internal syllables when no generalizations are required demonstrates that the relevant information can be encoded and discriminated, suggesting that participants' failure to generalize reflects something about the generalization process per se. Of course, we cannot rule out the possibility that future experiments could be devised to show successful symbolic generalization of sequence-internal repetitions. However, such a discovery would not explain the predominant patterns in our current data, and it would still be just as crucial to know how and why such small

---

[14] We are grateful to one of the anonymous reviewers for pointing out this possibility.

differences dramatically affect the ability to generalize the repetition-based structures used here, while having no effect on discrimination performance. The critical point here is that any successful generalization of sequence-internal repetitions, if possible, is at least much less robust than the ready generalizations in edge positions observed here, under neutral conditions, suggesting that—with a sufficiently sensitive experiment—fundamental differences in the capacity to generalize regularities in edge and nonedge positions can be observed. In other words, it seems unlikely that the advantage for edge positions may be explained by some sort of competence–performance distinction; rather, the advantage for edge positions in generalizing repetitions appears to be an intrinsic limitation of the underlying symbolic representational system.

Throughout all of these experiments, it is critical to note that the sequence edges were never prioritized over the sequence-medial positions in formal terms. As we have noted above, it can seem seductive to think that the edge positions were somehow intrinsically favored: Participants may have noticed whether a repetition was at an edge but only that it was "somewhere" in the middle of a sequence. This may be true in a sense, but if so it only serves to highlight the point we have tried to make here about the impact of perceptual salience on symbolic generalization. Note that in formal terms, subjects could just have readily noticed whether a repetition was at Positions 3 and 4 but only that it was somewhere near a sequence-final edge. The critical point here is that sequence-medial and sequence-final positions are distinguished by their relative salience, not by their formal structure: Although these positions seem intuitively different to us, they would be just as easily copied or shifted in a digital computer.

We conclude from these experiments, using repetition-based structures as a case study, that the spontaneous generalization of symbolic rules may not reflect a formal and fully general process such as that realized in a digital computer but may instead reflect the more piecemeal operation of *perceptual primitives*—more constrained symbolic operators that are triggered by the perceptual salience of only certain kinds of structures, based on Gestalt-like templates. In the rest of this article, we elaborate on this notion of perceptual primitives and then briefly discuss the implications of this framework and of our results for previous research and for other areas of cognitive processing such as language acquisition.

Perceptual Primitives

The notion of perceptual primitives to which we appeal differs from more global interpretations of symbolic processing in two primary ways, each related to the scope of the alleged symbolic operations. Perceptual primitives do not operate over all input patterns of a given type equally efficiently—that is, they do not operate over patterns without regard to the nature or content of those patterns. Rather, they are constrained both by the kinds of operations they can perform spontaneously and by the perceptual input conditions under which such operations are possible. In fact, Endress, Dehaene-Lambertz, and Mehler (2005) have recently provided some experimental support for the idea that repetitions may be special in this context: In a paradigm similar to that used in the present experiments, participants readily generalized repetition-based structures but failed to generalize other simple structures that did not involve repetitions. Using structures defined by pitch relations among piano tones, Endress et al. found that human adults readily learned the repetition-based structures ABA and ABB but not the ordinal structures low–high–middle and middle–high–low. This is especially striking given that associationist mechanisms are forced to predict that the ordinal structures should be processed more easily than the repetition-based structures. In any case, these results are consistent with the idea that repetitions are processed by a specialized symbolic mechanism.

In the present article, we have explored a complementary type of constraint on symbolic generalization, wherein even repetition-based structures themselves fail to be spontaneously generalized when they occur in certain sequence positions. We conclude that structures such as those based on repetitions may not be spontaneously generalized using the operation of a general mental symbol-manipulation system but rather by more limited operations that are constrained both by the kinds of available computations and by the perceptual input conditions under which they can operate. Such operations may constitute primitives for extracting structure from the environment in much the same way as certain key primitive structures and operations have been suggested to be crucial for visual object recognition (e.g., Biederman, 1987) or motor control (e.g., Fod, Matarić, & Jenkins, 2002; Krebs, Aisen, Volpe, & Hogan, 1999; Mussa-Ivaldi & Bizzi, 2000; Mussa-Ivaldi, Giszter, & Bizzi, 1994; Thoroughman & Shadmehr, 2000). In each of these views, difficult computational problems are solved using a divide-and-conquer strategy in which specialized efficient operators exist for perceiving or producing certain key structures, which can then be combined to process more complex structures and sequences. It is possible that the primitive at the core of our research may be a member of a fairly constrained inventory of such operators, each specialized for the extraction of particular types of structures. Under this interpretation, we would view participants' performance in our experiments as being driven by the use of a repetition-detection primitive that will only yield spontaneous generalization when "activated" with sufficient strength by salient positions such as edges.

Note that this notion of a perceptual primitive differs in a key way from previous proposals that have also appealed to computational primitives in symbolic processing (e.g., Block, 1995; Marcus, 2001; Pylyshyn, 1984). Such proposals were commonly motivated by an analogy between digital computers and the mind and have typically referred to the idea of *registers* and *instructions*. A central assumption of these proposals is typically that such instructions—even when highly constrained to implement a particular operation such as copying or comparing—still operate equally easily over all possible "values" of particular registers, simply because their operation does not depend on the registers' contents. In contrast, we suggest that such symbolic primitives are more constrained. Because they are sensitive only to particular salient patterns, they are constrained to operate only in certain configurations—for example, to detect edge repetitions but not sequence-medial repetitions. Such piecemeal primitives may constitute a more plausible (and conservative) view of symbol manipulation in our minds (compared with the more general operations in digital computers), if only because we already have some reason to believe that highly specialized and piecemeal computational abilities exist in many other organisms (e.g., Gallistel, 1990, 2000).

In this way, structure may be extracted from the environment using a collection of simple specialized (though possibly redundant) hardwired computational heuristics that may operate on different levels of representation, rather than using a monolithic formal symbolic structure-extraction process. We suggest that edge repetitions such as those studied here are members of this computational toolbox of the mind, and a longer-term goal of this project is to determine more generally which primitives may exist and how they are constrained. If this view is correct, it may be possible to investigate the underlying mechanisms (neural and otherwise) of particular symbolic operations, which would be rather difficult if one had instead to discover a single general basis for all computations that can be subsumed under the label *symbol manipulation*.[15] Of course, it will take much more evidence than is reported here to discover the truth of such an overarching proposal. Nevertheless, we suggest that the perceptual primitive view is a plausible interpretation that may serve as a baseline against which more general claims have to be evaluated.

In summary, previous interpretations have tended to treat repetitions as a case study for symbolic structure in general. In contrast, we suggest that not all types of symbolic relations may be equal: Repetitions may be efficiently detected and generalized (in edge positions) not because they constitute an example of symbolic structure in general but rather because they fit a particular primitive template for repetitions per se and occur in positions that are salient enough to trigger this template. We suggest that this perceptual primitive view is a plausible option that is consistent with both previous demonstrations of symbolic generalization in infants and adults and with our own data presented here, and that the distinction between perceptual primitives and general formal rule-extraction mechanisms is important (although it has not previously been made), because it may put constraints on possible theories of cognitive development.

Linguistic Implications?

Might perceptual primitives also play a role in language processing or acquisition? Although our stimuli and manipulations are clearly unlikely to have any direct implications for natural language, it is interesting to note that there are several types of evidence which suggest that perceptually salient positions can influence the acquisition of various aspects of grammar—including grammatical morphemes (e.g., Hsieh, Leonard, & Swanson, 1999; Johnston, 1991, as cited in Peters & Stömqvist, 1996; Peters & Stömqvist, 1996; Slobin, 1966, as cited in Peters & Stömqvist, 1996; Smoczynska, 1985, as cited in Peters & Stömqvist, 1996), the use of auxiliaries (e.g., Furrow, Nelson, & Benedict, 1979; Gleitman, Newport, & Gleitman, 1984; Newport, Gleitman, & Gleitman, 1977), and other syntactic constructions (e.g., Ingram & Thompson, 1996; Klein, 1974, as cited in Wijnen, Kempen, & Gillis, 2001; Wijnen et al., 2001).[16] Moreover, grammatical regularities such as the ban of repeated consonants in roots of Semitic languages (see McCarthy, 1985) are also subject to perceptual factors, because the probability of a violation of this regularity is a function of the similarity between the repeated consonants (e.g., Frisch, Pierrehumbert, & Broe, 2004). Edges also play an important role in theoretical linguistics, because many phonological and morphological regularities can be explained if the edges of constituents (e.g., syllables, feet, stems) are aligned (McCarthy & Prince, 1993).

Such results begin to suggest that perceptual primitives could play a role in some aspects of language. However, because nonhuman animals may also be endowed with such primitives, we have no reason to think that such operators are among the specific computational capacities that make language possible only in humans. Nevertheless, the language faculty could have recruited some perceptual primitives for its purposes, just as it may have recycled perceptual abilities present in other animals for the categorical perception of speech sounds (e.g., Kluender, Diehl, & Killeen, 1987), the compensation for coarticulation (Lotto, Kluender, & Holt, 1997), the sensitivity to rhythmical properties of different languages (Ramus, Hauser, Miller, Morris, & Mehler, 2000), or the sensitivity to transition probabilities among syllables (Hauser, Newport, & Aslin, 2001) and vowels (Newport, Hauser, Spaepen, & Aslin, 2004). In any case, it seems clear that linguistic structure involves many kinds of computational capacities, only some of which—if any—might be explained in terms of symbolic perceptual primitives.

Conclusions

The primary goal of this article has been to theoretically highlight and to experimentally address a type of ambiguity concerning the ability to spontaneously generalize symbolic regularities. Previous demonstrations of symbolic generalization have received considerable (and well-deserved) attention, in part because such abilities suggest that in a way the mind may operate in a manner similar to a digital computer. Our results suggest that such allusions may be appropriate in one sense but perhaps inappropriate in another sense. On the one hand, our results clearly support the view that such abilities reflect bona fide symbolic computation as involving type-based generalizations over repetition-based structures that are treated as variables. On the other hand, unlike symbolic operations in digital computers, our results suggest that such operations may not operate equally readily over all inputs. The edge-based repetitions that were spontaneously generalized in our experiments and in previous research may, after all, be a special case. Such abilities as in these experiments may not in fact serve as an example of a universal symbolic processing ability. Rather, such abilities may reflect the operation of perceptual primitive operations, for which edge-based repetitions are particularly salient. This type of piecemeal computation may suffice to

---

[15] This is related to a point made by Jerry Fodor (1983) in his seminal discussion of the "modularity of mind." He noted that cognitive science has been relatively successful in explaining certain encapsulated input systems in the mind—for example, aspects of vision and language—but has perhaps been less successful in explaining aspects of (less encapsulated) higher level cognition, resulting in Fodor's "First Law of the Nonexistence of Cognitive Science: . . . The more global . . . a cognitive process is, the less anybody understands it" (p. 107). Fodor's idea is that mental encapsulation allows for methodological encapsulation in the practice of cognitive science, by allowing for a successful divide-and-conquer strategy—whereas unencapsulated systems cannot be so readily divided. Perhaps a similar argument could be made here, where a piecemeal view of perceptual primitives may also allow for dividing and conquering.

[16] The influence of perceptual salience has also been observed for word recognition in infants, who recognize words better in utterance-final positions than in utterance-medial positions (Fernald, McRoberts, & Herrera, 1992; as cited in Fernald, Pinto, Swingley, Weinberg, & McRoberts, 1998).

explain symbolic generalization without thereby implicating any general monolithic view of cognitive processing. As in animal learning research, where careful study of the computational properties of particular processes has yielded important insights into the underlying (symbolic) mechanisms (e.g., Gallistel, 1990; Gallistel & Gibbon, 2002), a psychological theory of symbolic generalization can only be constructed by carefully studying the computational properties of specific operations. This requires not only demonstrating their flexibility but also exploring their constraints.

References

Altmann, G. T. (2002). Learning and development in neural networks—The importance of prior experience. *Cognition, 85,* B43–B50.

Altmann, G. T., & Dienes, Z. (1999, May 7). Rule learning by seven-month-old infants and neural networks. *Science,* 284, 875a.

Anderson, J. R. (1989). A theory of human knowledge. *Artificial Intelligence, 40,* 313–351.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Aslin, R. N., Saffran, J., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science, 9,* 321–324.

Bates, E., & Elman, J. L. (1996, December 13). Learning rediscovered. *Science, 274,* 1849–1850.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115–117.

Block, N. (1995). The mind as the software of the brain. In E. E. Smith & D. N. Osherson (Eds.), *An invitation to cognitive science* (2nd ed., Vol. 3, pp. 377–425). Cambridge, MA: MIT Press.

Chomsky, N. (1980). *Rules and representations.* Oxford, England: Blackwell.

Christiansen, M., Conway, C., & Curtin, S. (2000). A connectionist single-mechanism account of rule-like behavior in infancy. In L. R. Gleitman & A. K. Joshi (Eds.), *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 83–88). Mahwah, NJ: Erlbaum.

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

Dutoit, T., Pagel, V., Pierret, N., Bataille, F., & van der Vreken, O. (1996). The MBROLA project: Towards a set of high-quality speech synthesizers free of use for non-commercial purposes. In *Proceedings of the Fourth International Conference on Spoken Language Processing'96, 3,* 1393–1396.

Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology* (H. A. Ruger & C. E. Bussenius, Trans.). Retrieved June 18, 2005, from http://psychclassics.yorku.ca/Ebbinghaus/. (Original work published 1885)

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

Elman, J. L., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development.* Cambridge, MA: MIT Press.

Endress, A. D., Dehaene-Lambertz, G., & Mehler, J. (2005). *Perceptual constraints and the learnability of simple grammars*. Manuscript submitted for publication.

Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science, 6,* 205–254.

Fernald, A., Pinto, J. P., Swingley, D., Weinberg, A., & McRoberts, G. W. (1998). Rapid gains in speed of verbal processing by infants in the 2nd year. *Psychological Science, 9,* 228–231.

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science, 12,* 499–504.

Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences, USA, 99,* 15822–15826.

Fod, A., Matarić, M. J., & Jenkins, O. C. (2002). Automated derivation of primitives for movement classification. *Autonomous Robots, 12,* 39–54.

Fodor, J. A. (1975). *The language of thought.* Cambridge, MA: Harvard University Press.

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

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28*(1–2), 3–71.

Frisch, S. A., Pierrehumbert, J. B., & Broe, M. B. (2004). Similarity avoidance and the OCP. *Natural Language and Linguistic Theory, 22,* 179–228.

Furrow, D., Nelson, K., & Benedict, H. (1979). Mothers' speech to children and syntactic development: Some simple relationships. *Journal of Child Language, 6,* 423–442.

Gallistel, C. (1990). *The organization of learning.* Cambridge, MA: MIT Press.

Gallistel, C. (2000). The replacement of general-purpose learning models with adaptively specialized learning modules. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (2nd ed., pp. 1179–1191). Cambridge, MA: MIT Press.

Gallistel, C., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review, 107,* 289–344.

Gallistel, C., & Gibbon, J. (2002). *The symbolic foundations of conditioned behavior.* Mahwah, NJ: Erlbaum.

Gasser, M., & Colunga, E. (2000). Babies, variables, and relational correlations. *Annual Conference of the Cognitive Science Society, 22,* 160–165.

Gigerenzer, G., Todd, P., & The ABC Group. (1999). *Simple heuristics that make us smart.* New York: Oxford University Press.

Gleitman, L. R., Newport, E. L., & Gleitman, H. (1984). The current status of the motherese hypothesis. *Journal of Child Language, 11,* 43–79.

Gómez, R., & Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition, 70,* 109–135.

Gómez, R., & Gerken, L. (2000). Infant artificial language learning and language acquisition. *Trends in Cognitive Sciences, 4,* 178–186.

Gómez, R., Gerken, L., & Schvaneveldt, R. (2000). The basis of transfer in artificial grammar learning. *Memory & Cognition, 28,* 253–263.

Hauser, M. D., Newport, E. L., & Aslin, R. N. (2001). Segmentation of the speech stream in a non-human primate: Statistical learning in cotton-top tamarins. *Cognition, 78,* B53–B64.

Hinton, G. E., & Anderson, J. (Eds.). (1981). *Parallel models of associative memory.* Hillsdale, NJ: Erlbaum.

Hsieh, L., Leonard, L., & Swanson, L. (1999). Some differences between English plural noun inflections and third singular verb inflections in the input: The contributions of frequency, sentence position, and duration. *Journal of Child Language, 26,* 531–543.

Ingram, D., & Thompson, W. (1996). Early syntactic acquisition in German: Evidence for the modal hypothesis. *Language, 72,* 97–120.

Joanisse, M., & Seidenberg, M. (1999). Impairments in verb morphology after brain injury: A connectionist model. *Proceedings of the National Academy of Sciences, USA, 96,* 7592–7597.

Kluender, K., Diehl, R., & Killeen, P. (1987, September 4). Japanese quail can learn phonetic categories. *Science, 237,* 1195–1197.

Krebs, H., Aisen, M., Volpe, B., & Hogan, N. (1999). Quantization of continuous arm movements in humans with brain injury. *Proceedings of the National Academy of Sciences, USA, 96,* 4645–4649.

Lehman, J. F., Laird, J. E., & Rosenbloom, P. (1998). A gentle introduction to soar, an architecture for human cognition. In D. Scarborough & S. Sternberg (Eds.), *An invitation to cognitive science* (2nd ed., Vol. 4, pp. 211–253). Cambridge, MA: MIT Press.

Lotto, A., Kluender, K., & Holt, L. (1997). Perceptual compensation for coarticulation by Japanese quail (*Coturnix coturnix japonica*). *Journal of the Acoustical Society of America, 102*(2, Pt. 1), 1134–1140.

Marcus, G. F. (1998a). Can connectionism save constructivism? *Cognition, 66,* 153–182.

Marcus, G. F. (1998b). Rethinking eliminative connectionism. *Cognitive Psychology, 37,* 243–282.

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

Marcus, G. F. (1999b). Reply to Christiansen and Curtin. *Trends in Cognitive Sciences, 3,* 290–291.

Marcus, G. F. (1999c). Reply to Seidenberg and Elman. *Trends in Cognitive Sciences, 3,* 289.

Marcus, G. F. (1999d, April 16). Response to Seidenberg and Elman, Eimas and Negishi. *Science, 284,* 436–437.

Marcus, G. F. (1999e, May 7). Rule learning by seven-month-old infants and neural networks: Response to Altmann and Dienes. *Science, 284,* 875a.

Marcus, G. F. (2001). *The algebraic mind: Integrating connectionism and cognitive science.* Cambridge, MA: MIT Press.

Marcus, G. F., Vijayan, S., Rao, S. B., & Vishton, P. (1999, January 1). Rule learning by seven-month-old infants. *Science, 283,* 77–80.

McCarthy, J. J. (1985). *Formal problems in semitic phonology and morphology.* New York: Garland Press.

McCarthy, J. J., & Prince, A. (1993). Generalized alignment. In G. Booij & J. van Marle (Eds.), *Yearbook of morphology 1993* (pp. 79–153). Boston: Kluwer.

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

McClelland, J. L., Rumelhart, D. E., & The PDP Research Group. (Eds.). (1986). *Parallel distributed processing* (Vol. 2). Cambridge, MA: MIT Press.

Mussa-Ivaldi, F., & Bizzi, E. (2000). Motor learning through the combination of primitives. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences, 355,* 1755–1769.

Mussa-Ivaldi, F., Giszter, S., & Bizzi, E. (1994). Linear combinations of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences, USA, 91,* 7534–7538.

Negishi, M. (1999, April 16). Do infants learn grammar with algebra or statistics? *Science, 284,* 435.

Neurobehavioral Systems. (n.d.). Presentation (Version 0.47) [Computer software]. Albany, CA: Author.

Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4,* 135–183.

Newell, A., & Simon, H. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery, 19,* 113–116.

Newport, E. L., & Aslin, R. N. (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology, 48,* 127–162.

Newport, E. L., Gleitman, H., & Gleitman, L. R. (1977). Mother, I'd rather do it myself: Some effects and non-effects of maternal speech style. In C. Snow & C. A. Ferguson (Eds.), *Talking to children: Language input and interaction* (pp. 109–149). Cambridge, England: Cambridge University Press.

Newport, E. L., Hauser, M. D., Spaepen, G., & Aslin, R. N. (2004). Learning at a distance: II. Statistical learning of non-adjacent dependencies in a non-human primate. *Cognitive Psychology, 49,* 85–117.

Peña, M., Bonatti, L. L., Nespor, M., & Mehler, J. (2002, October 18). Signal-driven computations in speech processing. *Science, 298,* 604–607.

Peters, A., & Stömqvist, S. (1996). The role of prosody in the acquisition of grammatical morphemes. In J. L. Morgan & K. Demuth (Eds.), *Signal to syntax: Bootstrapping from speech to grammar in early acquisition* (pp. 215–232). Mahwah, NJ: Erlbaum.

Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition, 28*(1–2), 73–193.

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

Ramachandran, V. (1990). Interactions between motion, depth, color and form: The utilitarian theory of perception. In C. Blakemore (Ed.), *Vision: Coding and efficiency* (pp. 346–360). New York: Cambridge University Press.

Ramus, F., Hauser, M. D., Miller, C., Morris, D., & Mehler, J. (2000, April 14). Language discrimination by human newborns and by cotton-top tamarin monkeys. *Science, 288,* 349–351.

Rumelhart, D. E., McClelland, J. L., & The PDP Research Group. (Eds.). (1986). *Parallel distributed processing* (Vol. 1). Cambridge, MA: MIT Press.

Saffran, J., Aslin, R. N., & Newport, E. L. (1996, December 13). Statistical learning by 8-month-old infants. *Science, 274,* 1926–1928.

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

Saffran, J., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language, 35,* 606–621.

Seidenberg, M. (1997, March 14). Language acquisition and use: Learning and applying probabilistic constraints. *Science, 275,* 1599–1603.

Seidenberg, M., & Elman, J. (1999a, April 16). Do infants learn grammar with algebra or statistics? *Science, 284,* 433.

Seidenberg, M., & Elman, J. (1999b). Networks are not "hidden rules." *Trends in Cognitive Sciences, 3,* 288–289.

Shultz, T. R., & Bale, A. C. (2001). Neural network simulation of infant familiarization to artificial sentences: Rule-like behavior without explicit rules and variables. *Infancy, 2,* 501–536.

Thoroughman, K., & Shadmehr, R. (2000, October 12). Learning of action through adaptive combination of motor primitives. *Nature, 407,* 742–747.

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

Wijnen, F., Kempen, M., & Gillis, S. (2001). Root infinitives in Dutch early child language: An effect of input? *Journal of Child Language, 28,* 629–660.

Yang, C. D. (2004). Universal Grammar, statistics or both? *Trends in Cognitive Sciences, 8,* 451–456.

Received February 3, 2004
Revision received March 16, 2005
Accepted March 21, 2005 ■