The effect of repetition and similarity on sequence learning

Padraic Monaghan[1] and Chris Rowson[2]

[1]Lancaster University, UK

[2]University of York, UK

Word count: 3958

Abstract word count: 137

Corresponding author:

Padraic Monaghan

Department of Psychology

Lancaster University

Lancaster

LA1 4YF

UK

E-mail: p.monaghan@lancaster.ac.uk

Tel: +44 1524 593813

Fax: +44 1524 593744

Abstract

Repetition is a pervasive feature of children's environments, and may be an important contributor to learning such complex sequential structures as language. Endress, Dehaene-Lambertz, and Mehler (2007) found that repeated tone sequences were learned more easily than sequences containing ordinal relations, however, there have been no direct comparisons of repeating sequences versus sequences that contain similar, but non-identical, stimuli. In Experiment 1 we compared learning from repeating tone sequences to learning from tones that varied in similarity, and confirmed that repetition is a special case for learning. In Experiment 2 we showed that the learning distinction between repeated and similar elements is not affected by whether similarity is variable. We conclude by indicating that repetition provides an important constraint on learning, and discuss the extent to which such constraints are consistent with general-purpose statistical learning mechanisms.

Repetitions in stimuli have important consequences for perception and cognition (e.g., Henson, 1998; Kanwisher, 1987), and are thus likely to provide important constraints for learning. Repetition is a likely candidate as a useful feature for sequence learning, particularly for computing complex sequences such as those engaged for language learning (e.g., Saffran, 2003). Children's books and child-directed speech, for instance, are replete with instances of repeating words and phrases (Papousek, Papousek, & Haekel, 1987), which makes it seem unlikely to be an accidental property. Determining the precise role of repetition for learning sequence structure is therefore important in order to establish the identity of contributors to learning structure from sequences.

In a recent intriguing study, Endress, Dehaene-Lambertz, and Mehler (2007) examined the influence of repetitions in learning sequences of tones. They compared participants' learning of patterns of three musical tones that varied in terms of whether they contained repetitions or "ordinal" relations (in this case, whether the third tone was higher or lower in pitch than the first tone). Participants listened to four sets of three tones, then judged whether a fifth set of tones had the same or different structure. In the repetition condition, participants learned to distinguish low-high-low or high-low-high tones (ABA) from low-high-high or high-low-low (ABB) patterns, examples of these stimuli are shown in the top two rows of Figure 1. In the ordinal condition, the distinction was between low-high-medium (LHM) and medium-high-low (MHL) tones, third and fourth rows of Figure 1. Endress et al. (2007) found that participants performed more accurately in the repetition condition than the ordinal condition. However, there are several ways in which the repeating and the ordinal sequences vary in Endress et al.'s (2007) study, not only in terms of whether they contained a repetition, and we summarise these in the introduction to the first experiment. Our first aim for the current experiments was to provide a more rigorous test of the advantage of learning from sequences containing repetition by directly comparing repeating and non-repeating sequences.

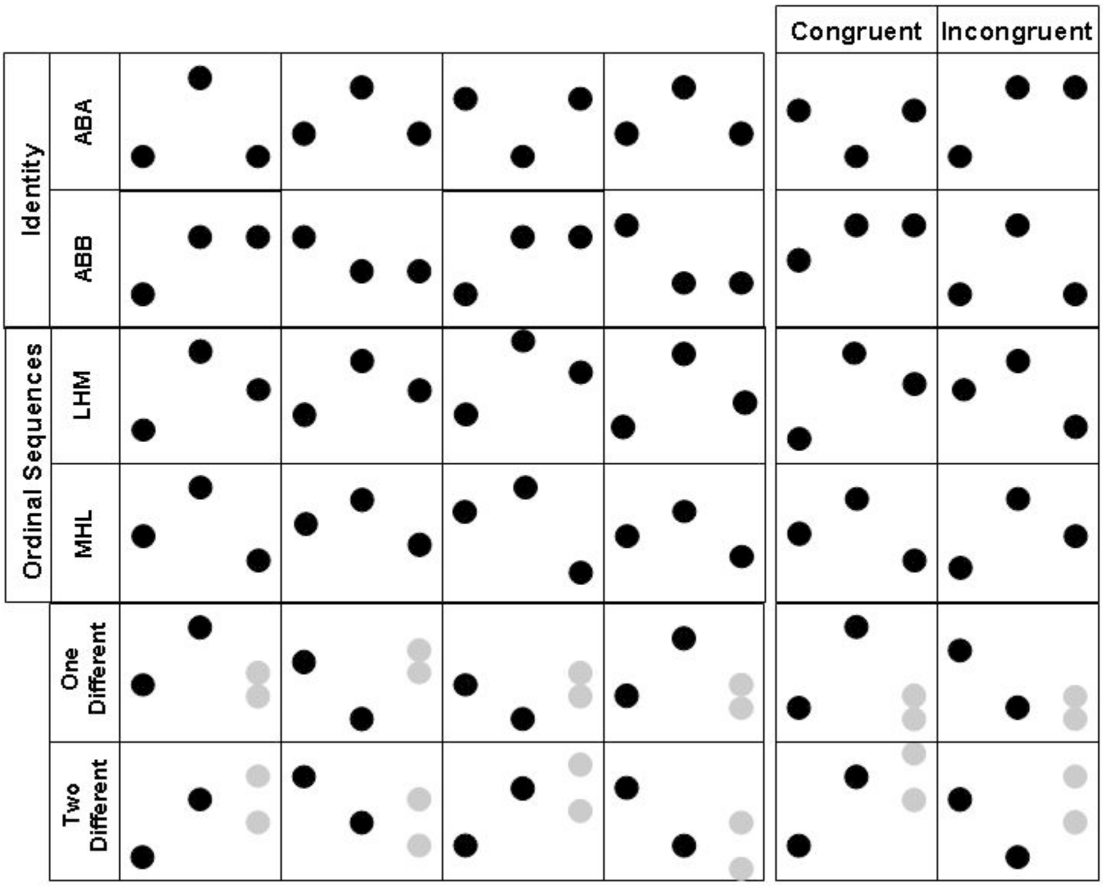

Figure 1. ABA and ABB tone patterns for repetition condition and MHL and LHM patterns from Endress et al. (2007), and ABA pattern for one-different and ABB pattern for two-different conditions in Experiment 1. Black circles indicate the pitch of the tone, and grey circles indicate the two possible pitches for the third tone.

A second aim of the current research was to determine whether repetition of identical stimuli is a special case for learning, or whether it is better conceived as one endpoint of a continuum of similarity. Previous studies of statistical learning in language-like sequences have provided evidence for similarity (rather than identical repetition) assisting in sequence learning. Newport and Aslin (2004) tested participants' ability to learn non-adjacent dependencies between two syllables when the intervening syllable was kept constant. They varied the extent to which the first and third syllables shared phonological properties, and found that learning was best when the first and third syllables shared phonological properties. Relatedly, Onnis et al. (2005) found that the non-adjacency structure of words in continuous speech was detected only when the dependent syllables shared phonological properties. Bonatti, Peña, Nespor, and Mehler (2005) also found learning of sequence structure when the dependent elements of speech-like sequences carried similarity in either the consonants or vowels.

Other studies have directly tested repetition in sequence learning, and found advantages for certain patterns of repetition in both infants (Marcus, Vijayan, Rao, & Vishton, 1999) and recurrent neural networks (Altmann, 2002). Yet, there has been no study that directly compares the effects of similarity to repetition. One possibility is that increasing similarity corresponds with increasing accuracy of learning, so the effect of repetition is quantitatively continuous with sequences containing similar stimuli. Another possibility is that repetition is instead qualitatively distinct from sequences, so the distinction in learning from sequences containing similar stimuli would be a step change from sequences containing precise repetitions. Establishing this aspect of repetition processing has implications for the constraints required in an adequate cognitive model of sequence learning. If repetition is continuous with similarity then this conforms with standard views of statistical learning mechanisms applied to sequences, where the similarity of representation of stimuli is related to their featural similarity – coherent with one of the foundational principles of connectionist models (Hopfield, 1982; Kohonen, 1988). However, if repetition is qualitatively distinct from similarity, then this presents a challenge to models of statistical learning that represent stimuli along a continuum according to similarity.

A further aim of the current studies was to tease apart the contribution of similarity between two items, and the contribution of variability in how predictable the similar elements actually are. Sequences that contain similar, but non-identical, stimuli may be predictably different, as in the Newport and Aslin (2004) and Onnis et al. (2005) studies –where the stimuli with similar properties are reliably co-occurrent. Alternatively, similar sequence stimuli may differ in their predictability – where the extent of the similarity is variable. Such differences in predictability, may be an important contributor to potential distinctions between repeating and non-repeating sequences.

The first experiment compares learning of sequences where either the first and third or second and third tones differ by zero, one, or two semitones, though the final tone can be either higher or lower. This experiment examines the combined influence of similarity and variability of the solution space. The second experiment tests only the effect of similarity, where the third tone was either identical or higher by one or two tones from the first or second tone.

Experiment 1: Learning from similarity and variability

We used a paradigm related to Endress et al.'s (2007) study to determine the role of repetition in sequence learning. However, the repeated versus ordinal sequence comparison in Endress et al.'s study introduced differences in addition to whether the sequence contained repetitions or not. First, ABA and ABB sequences could be differentiated just by attending to the final two tones in each sequence and responding whether they are the same or different. For ordinal sequences, both MHL and LHM sequences can only be contrasted by reference to the first and third tone (note that participants were trained on either the repetition or the ordinal sequences). Such non-adjacencies are more difficult to learn than adjacent dependencies (Gómez, 2002; Onnis, Monaghan, Richmond, & Chater, 2005; see also indirect evidence in Peña, Bonatti, Nespor, & Mehler, 2002). Relatedly, the non-adjacent dependencies in the ordinal sequences were likely to be difficult to identify because of the lack of variation in the second tone – this is always a tone of higher pitch. In the repeating sequences, the second tone may rise or fall with respect to the first tone. Gómez (2002) found that increasing variability in the intervening element in three-word sequences improved learning of the relationship between the first and the third word. Lack of variability promotes bigram monitoring (Onnis, Christiansen, Chater, & Gómez, 2003) – an ineffective strategy for the ordinal sequences.

Second, the search space of the final tone was greater in the ordinal sequences than the repetition sequences. In the ordinal condition, the final tone of each sequence can be anywhere below the first tone in MHL sequences, and anywhere between the first two tones in the LHM sequences. In the repetition sequences, the final tone is the same as either the first or the second tone, and so the potential search space is smaller and more reliably defined. The third experiment in Endress et al.'s paper addressed this unpredictability in ordinal sequences by maintaining the interval between the first and third tone a constant two tones, and this improved learning.

We adapted the ABA/ABB sequences from Endress et al. (2007) to test whether identical repetition was advantageous in learning over conditions where the final tone varied from the first or second tone by either one or two semitones. If similarity contributes in a graded manner to learning, then the identical repetition condition should be learned more easily than the one-semitone condition, which in turn should be learned more easily than the two-semitone condition. If identical repetition is qualitatively distinct from similarity, then only the identical repetition condition should differ from the one- and two-semitone condition. We here assume that differences of zero, one and two semitones are judged by participants to be declining in similarity. Harris (1985) directly tested participants' judgements of tone similarity by playing pairs of tones, and found that tones distinct by two semitones were judged to be less similar than one semitone distinct tones.

Method

Participants. Nine male and 21 female students at the University of York participated in the experiment for course credit or for £2.

Materials. 20 piano monotones in semitone intervals were generated using a synthesiser (Nyquist, 2007). Semitones ranged from G (97 Hz) to a second E-flat(313 Hz). Each semitone was sampled at 22,050Hz with 16 bit resolution, and lasted for 400 ms.

There were three conditions: “Identical repetition”, “one-different”, and “two-different” (see Figure 1). For ABA and ABB patterns in the identical repetition condition, the second tone was either higher or lower than the first by a minimum of six semitones, and the third tone was identical to the first tone in the ABA pattern, and identical to the second tone in the ABB condition. In the one-different condition, the third tone was one semitone higher or lower than the first tone in the ABB condition, and ABB patterns were the same except third tones were one semitone higher or lower than the second tone. Tones in the two-different condition were two semitones higher or lower than the other tone.

Procedure. Each trial consisted of five piano tone triplets, played through closed cup headphones. The number of each triplet presented was shown on the screen (1,2,3,4) and when “?” appeared, participants were required to decide whether the last triplet conformed to the others (by pressing “y” on a keyboard), or was a different pattern (pressing “n”). Congruent and incongruent trials were presented randomly, and feedback (correct/incorrect) was presented on the screen after each trial. There were 160 trials in total, of which half were congruent, and the experiment lasted roughly 30 minutes.

Results and Discussion

The 160 responses were broken down into 8 blocks of 20 responses each. Accuracy for each block was entered as the dependent variable in an ANOVA with time (blocks 1 to 8), experimental condition (identity/one-different/two-different), and sequence type (ABA/ABB) as factors.

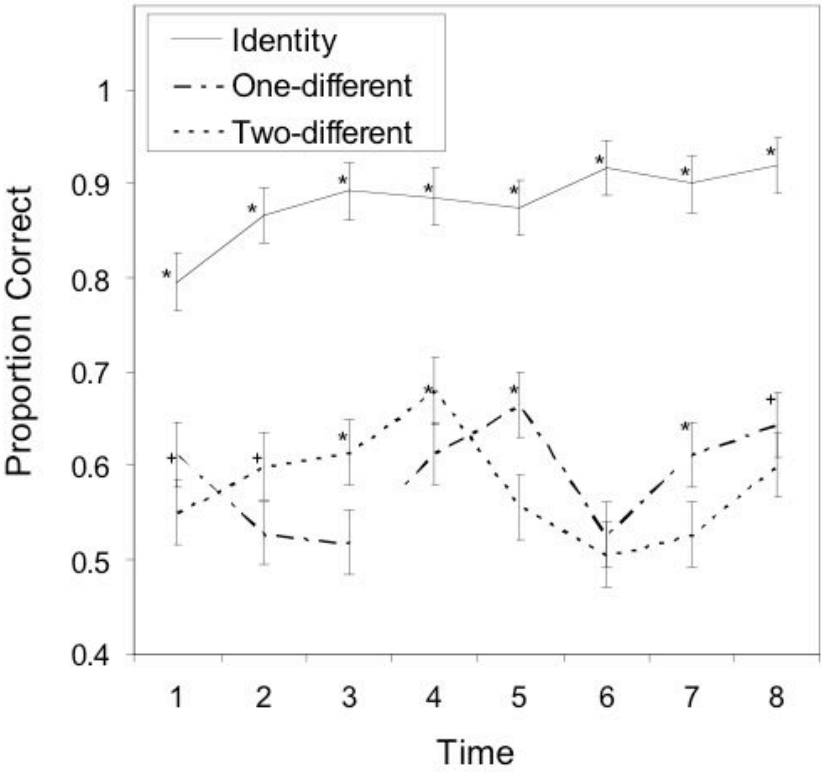

Figure 2. Estimated marginal means for Experiment 1, with variability for one- and two-semitone-different conditions. Indices next to data points indicate difference from chance level, * p < .05, + p < .1.

There was a main effect of time, *F*(7, 189) = 2.50, $p < .05$, $\eta^2 = .09$, with accuracy increasing with time. There was a significant main effect of condition, *F*(1, 27) = 50.18, $p < .001$, $\eta^2 = .79$, with learning significantly higher in the identical repetition condition than the one-different and two-different conditions, Tukey's test $p < .001$. The one-different and two-different conditions were not significantly different, $p = .68$. The interaction between time and condition was also significant, $F(14, 189) = 2.61$, $p < .005$, $\eta^2 = .16$, indicating that there was greater improvement in the identical repetition condition than in the other conditions. No other main effects or interactions were significant, all $F < 1$. Comparisons with chance-level performance are shown at each time step in Figure 2. Over all time steps, the identical repetition condition and the two-different condition resulted in significantly better learning than chance, $p < .01$, but the one-different condition was only marginally better than chance after Bonferroni correction $p = .07$, A summary of the main effect of condition is shown in Figure 5a. Harris (1985) found that the difference in similarity judgements between one and two semitone different pairs of tones was smaller than that between zero and one semitone different pairs, and so the sequence learning results match the psychoacoustic properties of tone similarity.

The results supported the claim made in Endress et al. (2007) that identical repetition is a special case for learning. Whether the third tone differed by one or two semitones from the first or the second tone had a significant and equal impact on learning compared to when there was no difference in pitch for the third tone. However, the one- and two-different conditions differed from the identical repetition condition in terms of whether the third tone was precisely predictable or not from the first two tones as the tone could either rise or fall in pitch – so Experiment 1 tested the combined effects of increasing similarity and variability. Experiment 2 tested whether this effect of variability was critical for the results by repeating the study with the one- and two-different conditions always producing a rise in the third tone, instead of the tone either rising or falling.

Experiment 2: Learning from graded repetitions without variability

Method

Participants. 10 male and 20 female students from the University of York participated for course credit or a sum of £2. None had participated in Experiment 1.

Materials. The materials were identical to those of Experiment 1 except that the third tone in the one-different condition always rose by one semitone over the first tone for ABA sequences or over the second tone for ABB sequences. The two-different condition always rose by two semitones over the first tone for ABA sequences or the second tone for ABB sequences.

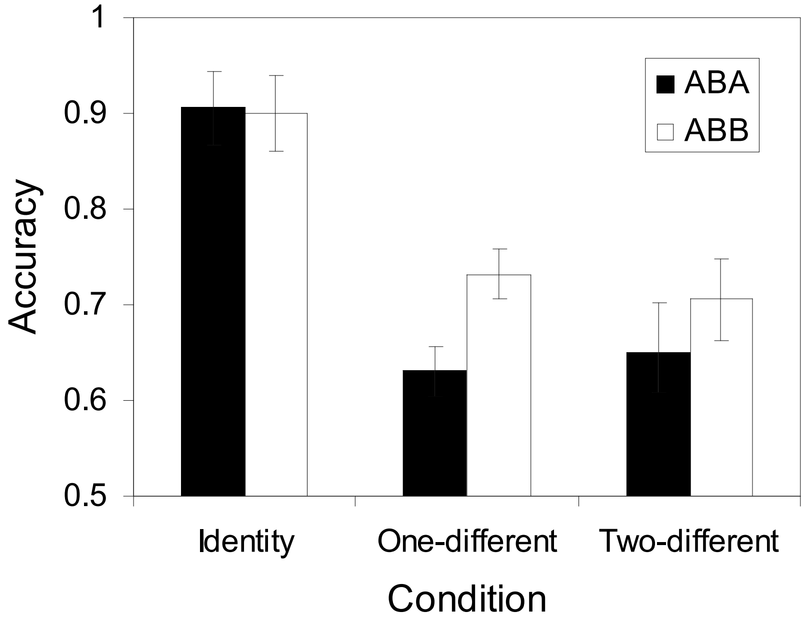

Figure 3. Effect of Condition on ABA and ABB sequences in Experiment 2.

Procedure. The procedure was identical to that of Experiment 1.

Results and Discussion

As for Experiment 1, we blocked the trials into groups of twenty. Time, condition, and sequence type were entered as factors in an ANOVA with accuracy as dependent variable. The main effects were qualitatively similar to those for Experiment 1: For time, $F(7, 189) = 2.90, p < .01, \eta^2 = .10$, and condition: $F(1, 27) = 12.57, p < .001, \eta^2 = .48$. Identical repetition

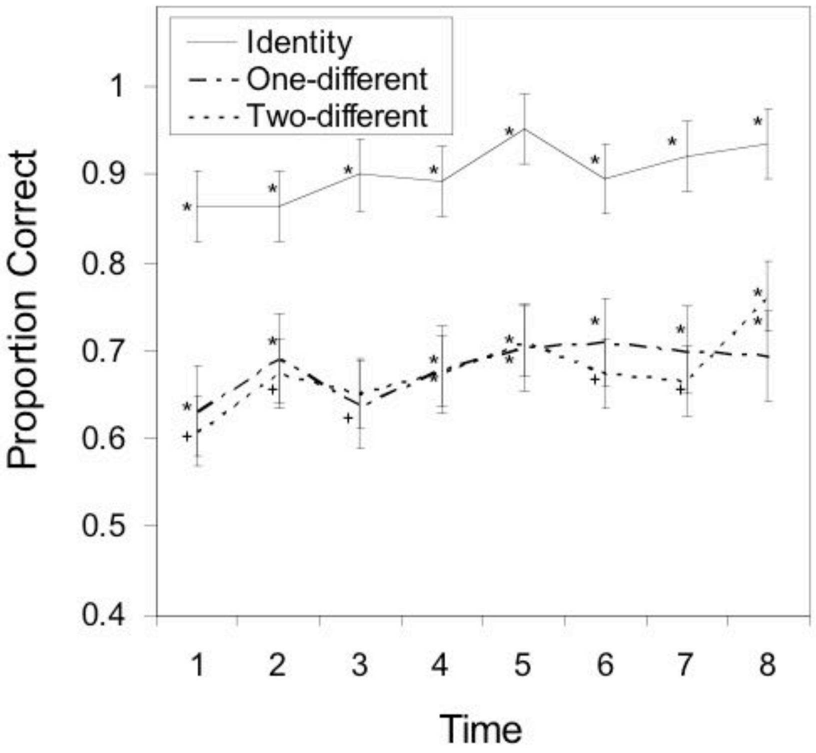

Figure 4. Estimated marginal means for Experiment 2, with no variability for one- and two-semitone-different conditions.

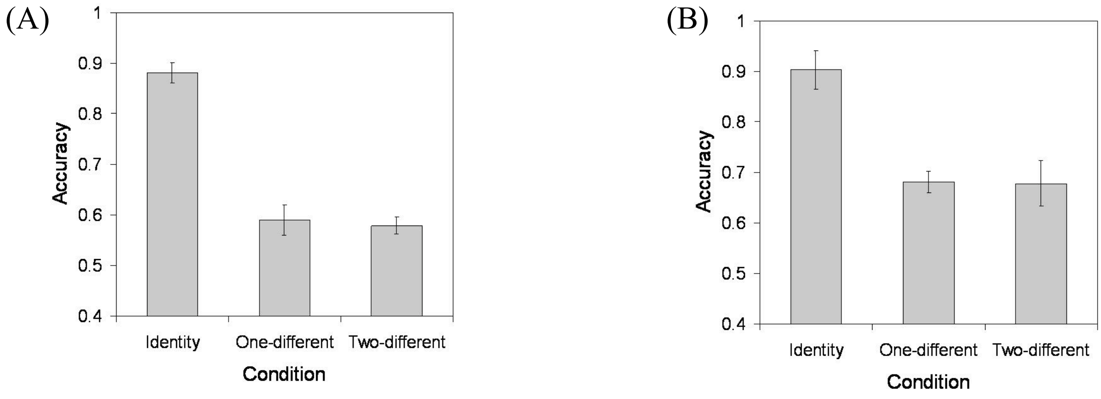

Figure 5. Summary of results from (A) Experiment 1 and (B) Experiment 2 showing means and standard error of the means across the three conditions.

resulted in significantly higher accuracy than the other conditions, both $p < .001$, but the one- and two-different conditions did not differ, $p = .99$. As in Experiment 1, the identical repetition condition was qualitatively distinct from the one- and two-different conditions, which resulted in a similar level of learning.

Distinct from Experiment 1, however, was a significant main effect of sequence type, $F(1, 27) = 11.69$, $p < .005$, $\eta^2 = .30$, and a significant interaction between condition and sequence type, $F(2, 27) = 4.56$, $p < .05$, $\eta^2 = .25$. This was due to no significant difference between ABA and ABB sequences for the identical repetition and two-different conditions, both $p > .2$, but a significant advantage for ABB over ABA sequences for the one-different condition, mean correct = .63 and .73, respectively, $p < .05$, see Figure 3. Predictable differences in pitch between tones that were very similar but not identical resulted in a slightly different strategy for learning the sequences, where perhaps the final bigram of each tone triple determined performance – participants listened for whether the final two tones rose by a single semitone, or whether they described a larger interval. The other interactions were not significant, all $F <= 1.09$. Comparisons to chance level for each time block are shown in Figure 4, Over all blocks, all three conditions were significantly better than chance, all $p < .01$, indicating that predictability resulted in less variable responses in the one- and two-different conditions, compared to Experiment 1. Figure 5b shows a summary of the main effect of condition.

General Discussion

The results of the two experiments met the first aim of the study, confirming that repetition has a special status in sequence learning. When tones were identical between first and third or second and third positions then learning was significantly better than when the relationship between first and third or second and third tones differed by one or two semitones. The second aim of the study was to determine whether precise repetition figured along a continuum of similarity, or whether it was qualitatively distinct from one- or two-semitone different sequences. We found that deviating from identical repetition by either a semitone or a tone had an equal and catastrophic impact on learning, indicating that, at least for tone sequences, there was a step-change in learning from repeated compared to non-repeated but similar sequences. This close experimental comparison between conditions of repetition and non-identical similarity goes beyond previous demonstrations of the importance of similarity for scaffolding learning of word or grammatical structure (Bonatti et al., 2005; Newport & Aslin, 2004; Onnis et al., 2005; Pothos & Bailey, 2000). in that we have shown that precise repetitions provide the strongest constraints on sequence learning.

The third aim was to determine the relative effect of variability on learning sequences with different degrees of similarity between tones. We found that variability did not have a profound quantitative impact on learning over and above the extent of similarity. For both Experiment 1 (with variability) and Experiment 2 (without) we found similar levels of performance for the one and two semitone different conditions. An ANOVA with experiment, condition, time, and sequence type revealed no significant main effect or interactions of experiment, all $F < 1$. However, there was indirect evidence that no variability resulted in better learning over time of the sequence structure – the significant interaction between time and condition from Experiment 1 (where the identical repetition resulted in faster learning) was not replicated in Experiment 2, because of some evidence of learning in each of the conditions in this study (compare Figures 2 and 4). Furthermore, the learning of ABA and ABB sequences appeared to be slightly different in the two experiments. In Experiment 1, there was no advantage of one sequence type over the other, however, in Experiment 2, for the one-different condition, ABB sequences were learned more easily than ABA sequences. This indicates that participants' performance may be somewhat strategic for these tasks – altering one characteristic of the stimuli, from variable to non-variable, resulted in different response to the stimuli. The better learning of ABB sequences, consistent with Endress et al.'s (2007) studies, was only reliably found for the one-semitone different condition, perhaps indicating sensitivity to bigrams for this contrast – participants preferentially respond to a rise by one semitone (in the ABB sequences) compared to a change greater than two semitones (as in the ABA sequence). However, though participants may be learning from bigrams in this condition, note that bigrams are an equally effective strategy in all conditions in the experiments. Thus the comparison between zero, one and two semitone different sequences can still be made directly, even if it is unclear whether participants are learning only from adjacent tones, or learn across non-adjacencies as well.

Whatever the locus of the effect of repetition in the statistical structure of the sequences, our studies have confirmed that repetition can guide learning, and that it is a special case as a constraint for detecting structures. Such effects are consistent with the effects of repetition in perceptual and memory tasks (e.g., Henson, 1998; Kanwisher, 1987), consequently repetition in speech to children is likely to have a profound effect in learning word boundaries and grammatical structure (e.g., Onnis et al., 2005). So, what sort of learning mechanism can account for these effects? An adequate learning system must be able to account for (1) sensitivity to repetition which is discontinuous with respect to similarity, and (2) learning of the structure that is guided by repetition. The current data does not constrain the architecture of the cognitive system for the latter point –repetition, once detected, may then pass onto either an associative learning mechanism (Altmann, 2002) or a symbolic processing system (Marcus et al., 1999).

So the point at issue is the implementation of repetition sensitivity in the cognitive system. Endress et al. (2007) claimed that the special status of repetition requires a modular perceptual mechanism that detects identity in sequences. But is repetition sensitivity also compatible with general-purpose learning mechanisms? The effects of similarity in sequences are likely to be available to computational models that respond to the statistical structure of the input. Indeed, detecting similarity (and representing such similarity within the model) is a fundamental principle in self-organising neural network models of cognition (e.g., Hopfield, 1982; Kohonen, 1988). Yet, such models would predict that repetition is on the same continuum as stimuli that are non-identical but perceptually similar, which was not the case in the current experiments. However, a statistically-based model learning sequences which contain repetitions would be required to distinguish different occurrences of the repeating stimuli at different points in the sequence, and such a mechanism could be generated without requiring a perceptual mechanism. One possible modeling framework, for instance, is the recurrent backpropagation model of short term memory by Botvinick and Plaut (2006), where the representation of stimuli in the hidden units of the model is a function of the stimulus identity and its sequence position. Whether such a model would reflect the qualitative distinction between identical repetition and similarity is an open question, but the respecting of both featural and temporal similarity in this model provides some promise for simulations of the human data without requiring a specific cognitive architecture for processing repetitions.

References

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

Bonatti, L.L., Peña, M., Nespor, M., & Mehler, J. (2005). Linguistic constraints on statistical computations: The role of consonants and vowels in continuous speech processing. *Psychological Science, 16*, 451-459.

Botvinick, M. & Plaut, D. C. (2006). Short-term memory for serial order: A recurrent neural network model. *Psychological Review, 113*, 201-233.

Endress, A.D., Dehaene-Lambertz, G., & Mehler, J. (2007). Perceptual constraints and the learnability of simple grammars. *Cognition*, *105*, 577-614.

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

Harris, R.W. (1985). Perceived relatedness of musical tones in major and minor tonal contexts. *The American Journal of Psychology, 98*, 605-623.

Henson, R.N.A. (1998). Item repetition in short-term memory: Ranschburg repeated. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 1162-1181.

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, 79*, 2554-2558.

Kanwisher, N.G. (1987). Repetition Blindness: Type recognition without token individuation. *Cognition, 27*, 117-143.

Kohonen, T. (1988). The "neural" phonetic typewriter. *Computer, 21*, 11-22.

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

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

Nyquist (http://www.cs.cmu.edu/~music/nyquist/) . Retrieved 6th May 2007.

Onnis, L., Christiansen, M.H., Chater, N., & Gómez, R. (2003). Reduction of uncertainty in human sequential learning: Evidence from artificial grammar learning. *Proceedings of the 25th Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Lawrence Erlbaum Associates.

Onnis, L., Monaghan, P., Richmond, K., & Chater, N. (2005). Phonology impacts segmentation in speech processing. *Journal of Memory and Language, 53*, 225-237.

Papousek, M., Papousek, H., & Haekel, M. (1987). Didactic adjustments in fathers' and mothers' speech to their 3-month-old infants. *Journal of Psycholinguistic Research, 16*, 491-516.

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

Pothos, E.M. & Bailey, T.M. (2000). The importance of similarity in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 26*, 847-862.

Saffran, J.R. (2003). Statistical Language Learning: Mechanisms and Constraints. *Current Directions in Psychological Science*, *12*, 110-114.