Chapter 17

Communicating structure, affect, and movement

Zoltán Dienes, Gustav Kuhn, Xiuyan Guo, and Catherine Jones

In Chapter 16, Bharucha, Curtis, and Paroo propose that music serves to communicate affect, the experience of motion, and an inducement to a particular structural interpretation. We will take the meaning of 'communicate' to be broad and not necessarily implying all the pragmatic constraints that successful communication often entails. For example, we will take a structure to be successfully communicated even if the recipient is not consciously aware of what the structure as such really is. Commonly people appreciate musical structure, and have the experiences intended by the composer, yet are not consciously aware of what the structure is. Indeed, this commentary will focus on the case where a structure put into music by us is detected by listeners without them being able to say what it is exactly they have detected. We show how some musical structures are analogous to certain linguistic structures in virtue of exhibiting mirror symmetries. We argue that people can come to implicitly learn to detect symmetries, musical inversions in particular (Dienes & Kuhn, forthcoming). Such implicit learning leads to greater liking of the structures learnt. Then we will show that just as music may communicate affect, structure, and movement, so can movement communicate structure and affect. Just as melody is a type of movement in tonal space, so can physical movement embody the same symmetry patterns of our melodies. Likewise, we present evidence that people can implicitly learn symmetries in movement. Finally we will discuss computational models of acquiring sensitivity to these structures in music and movement. First we will start with a brief introduction to implicit learning.

Implicit learning

The term 'implicit learning' was coined by Reber in 1967 to refer to the process by which we acquire unconscious knowledge of the structure of our environment. A key example that inspired Reber was natural language: Young children acquire knowledge of their native syntax in a way that does not allow them to verbalize what has been learned, even though they may have learned it very well indeed. Reber sought to investigate such learning in the lab using artificial grammars with adults. Whether or not there is a general purpose implicit learning process that applies in the same way to any structured domain, as Reber thought, he nonetheless could observe implicit learning—the acquisition of unconscious knowledge—using artificial grammars in the lab. He used an artificial finite state grammar to specify allowable strings of letters, such that the letter strings were structured but on casual inspection did not appear to be so. People were asked to memorize such strings of grammatical letters without being told there was a grammar. After a few minutes of exposure to the training strings, people were then informed that there was a complex set of rules determining the order of letters within strings. People were asked to classify new strings as being either grammatical or not. Reber found people could classify above chance (e.g. 65%) even though they had poor ability to verbalize relevant rules and often reported they were sorry to have messed up the experiment. On the face of it, people had acquired unconscious knowledge of the finite state grammar. Reber's work was ignored for a couple of decades but then triggered an extensive literature, debating what exactly it is that people learn in the artificial grammar learning paradigm (e.g. Pothos, 2007) and whether or not it is unconscious (e.g. Shanks, 2005). In terms of the latter issue, simply asking people to describe the rules, as Reber had originally done, may be an insensitive measure of conscious knowledge. Research subsequent to Reber's original papers has used more sensitive measures. The evidence for unconscious knowledge depends on which theoretical framework is used for defining consciousness (see Seth, Dienes, Cleeremans, Overgaard, & Pessoa, 2008, for discussion). If knowledge is conscious when one is conscious of having the knowledge, then there is good evidence that people do indeed acquire unconscious knowledge in Reber's artificial grammar learning paradigm (Dienes, 2008; cf. Shanks, 2005).

In terms of what is learned, Reber's original claim was that people had learned abstract knowledge of the finite state grammar. This was a good first guess because people could generalize from the training strings to novel strings generated by the grammar. Brooks (1978) made the first and most radical challenge to Reber's claim. Brooks pointed out that people could simply memorize whole training strings and classify a test string on the basis of its similarity to one or more training strings: no abstract knowledge need be induced in training at all. In general, according to this approach, conceptual knowledge could be based on storing all encountered exemplars of a concept, with no further attempt at abstraction. Computational versions of such exemplar models (also called 'case-based reasoning' in other contexts) have been tested in artificial grammar learning (e.g. Dienes, 1992; Jamieson and Mewhort, 2009) and can fit a range of findings, though not all (e.g. Pothos & Bailley, 2000). It is unlikely people implicitly learn about musical structure through memorizing whole musical pieces, but rather, if exemplar learning is involved, of fragments. Fragments of what size? In artificial grammar learning, people appear to implicitly learn fragments of two, three, and sometimes four letters (e.g. Servan-Schreiber & Anderson, 1990; Perruchet & Pacteau, 1990), and Wiggins (see commentary, Chapter 18, this volume) has modelled the learning of musical structure with similar sized fragments of musical notes.

As people can implicitly learn more than fragments, as we will discuss later, another approach to modelling the implicit learning of structure is to use connectionist networks, which can flexibly learn a range of structures. Connectionist networks can learn allowable fragments yet also produce representations which lie somewhere on a continuum of abstractness from exemplar models to rules of a grammar per se (Cleeremans, 1993; Dienes, 1992). Cleereman's used a simple recurrent network (SRN) and showed that it learned not only the conditional probabilities of successors to sequences of increasing length (i.e. fragment information) but also organized its internal representations in a way similar to the structure of a finite state grammar. The SRN was also adapted to learning the constraints in musical compositions by Mozer (1994). We will consider a specific abstract structure, mirror symmetry in music, and its learnability by the SRN below.

Bharucha et al. used a connectionist network to model musical knowledge consisting of nodes coding notes connected to nodes coding chords connected to nodes coding keys. One version of the model self-organized into the required hierarchical structure. Self-organizing maps have yet to be systematically explored in implicit learning generally, and this is a promising line of research. However, as we will see, Bharcha's model will need augmenting to deal with the symmetry structures we will consider in this paper.

Music communicates affect and structure

Zizak and Reber (2004) proposed that implicit learning of the structure of a domain can lead to greater preference for items having that structure: a structural mere exposure effect. They reviewed and presented new evidence that implicit learning of visual sequences obeying a finite-state grammar could lead to enhanced liking of novel sequences obeying the same grammar. That is, the communication of structure is the means by which affect is itself communicated. In the finite-state grammars used by Zizak and Reber, most of the structure in the grammar is accounted for by the ability of sequence elements to predict their immediate successors, i.e. by commonly occurring chunks of two or three sequence elements. Such chunking is an important part of learning music as well as language, but it does not account for all the structure that people can learn in either. The grammars above finite-state in the Chomskian hierarchy allow recursion and hence different types of symmetries. Figure 17.1 shows such two types of symmetry.

Consider a sequence of four elements, which we will arbitrarily call 'k l m n'. The letters of the alphabet are used simply to indicate the temporal order of the element (k comes before l etc.), but the identities of the elements are free to vary. Each element could, for example, be a tone, and the sequence a melody. We will label the notes of a particular melody: k1 l1 m1 n1. On the right hand side of Figure 17.1, this melody is indicated schematically. In the top it is reflected through a horizontal mirror. This creates a musical inversion. The diagram on the left indicates the mapping relation between notes (or intervals): the first note in the melody predicts the first note in the inversion; the second note in the melody predicts the second in the inversion and so on. The case of an inversion is illustrated in Figure 17.2 as well, showing that the mapping of intervals involves going to their opposite: if the melody initially descends three steps, the inversion initially ascends three steps and so on. Now consider the symmetry in the bottom part of Figure 17.1. Now the melody is reflected in a vertical mirror and a retrograde results. The mapping relation is shown on the left.

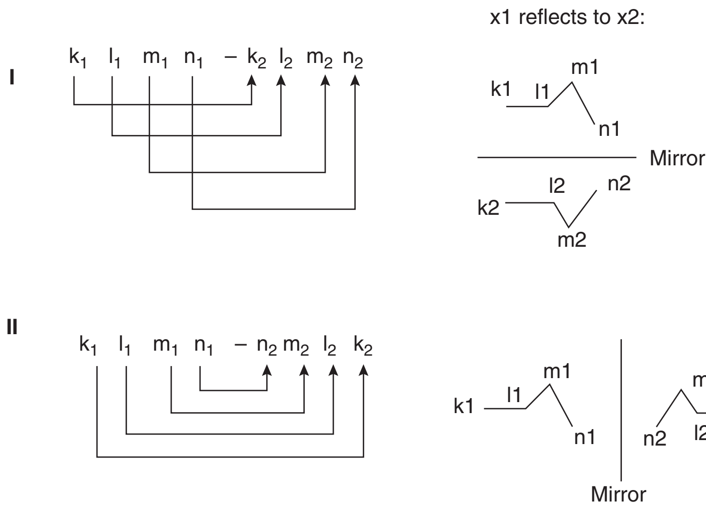

Fig. 17.1 Two types of symmetry.

The first mapping relation in Figure 17.1 corresponds to inversions and transposes, and also to certain linguistic structures. The 'respectively' construction has the same structure, for example: 'Tenzin, Trinley, and Tumpo wore yellow, black, and red hats, respectively', in which the first name goes with the first colour, the second with the second, and so on. The same structure is seen in cross serial dependencies in Dutch and Swiss German. Such dependencies require a grammar more powerful than context free. The second mapping relation in Figure 17.1 corresponds to musical retrogrades, and also to nesting structures in language like centre embedding: 'The bamboo the panda ate was fresh'. The first noun goes with the last verb and the last noun with the first verb. Such dependencies can be produced by context free grammars.

Most models of implicit learning were designed to account for the way people can implicitly learn chunks (for a review see e.g. Pothos, 2007); but no amount of chunk learning can make one sensitive to symmetries per se. A different sort of computational model is needed for detecting symmetries than forming chunks. Yet symmetries in an abstract way are part of language and a part of music, and also of movement: Most movements we can do, we can also do backwards to get back to where we started;

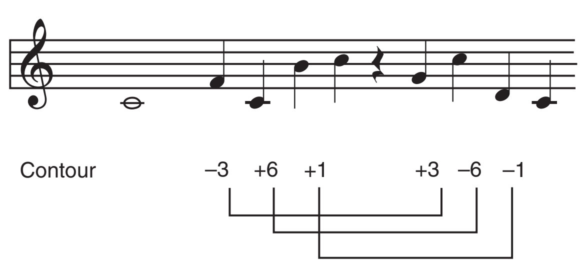

Fig. 17.2 A musical inversion.

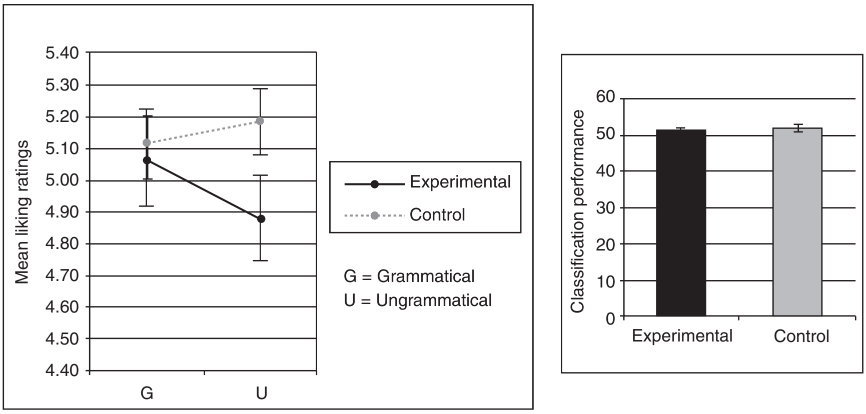

Fig. 17.3 Learning inversions, Kuhn and Dienes (2005).

and movements we do with the left hand we can do mirror reflected with the right, while appreciating they are in some sense the same. It would make evolutionary sense for perceptual systems to learn to detect symmetries, if only because detecting symmetries allows shorter length encodings of stimuli and hence easier storage. Are mirror symmetries indeed the sort of structures that we can implicitly learn and that can thus influence our affect?

Kuhn and Dienes (2005) asked participants to listen to 120 tunes, each tune consisting of eight notes. The last four notes were always an inversion of the first four, though participants were not informed of this relation, nor indeed that there were any structural invariants in the melodies: participants just tried to detect for each melody whether it had occurred before, a cover task which served simply to keep them concentrating on each tune as a whole. Then subjects either rated how much they liked a set of new strings or classified them. The new test strings were composed of novel pairings of tunes and of intervals, so the chunks of which they were constructed had never occurred in the training phase. Half the new trains had the same inversion structure as the training tunes, and half did not. For the classification task subjects were told that the strings they had been listening to were generated by a rule and they were about to hear some new strings, half of which followed the rule. They were asked to classify half of the tunes as following the rule and half as not. We also ran a control group that had no training but just completed the test phase. As shown in Figure 17.3, while the control subjects preferred the non-inversions to the inversions, the trained subjects preferred inversions. Exposure to inversions led to a relative liking of inversions. Knowledge that the rule in question was inversion appears to be unconscious: trained subjects could not classify the strings as rule following or not at above chance levels. Melody communicated structure and affect (see also Dienes & Kuhn, submitted).

Dienes and Longuet-Higgins (2004) used more complex sequences that followed the constraints of serialist music, i.e. each sequence was a 'tone-row' (i.e. 12 notes long and each note A to G occurred once and only once). People were asked to assess whether the 'reply' (second half of the tone row) went with the 'theme' (first half). The tone rows could instantiate serialist transposes, inverses, retrogrades, or inverse retrogrades.

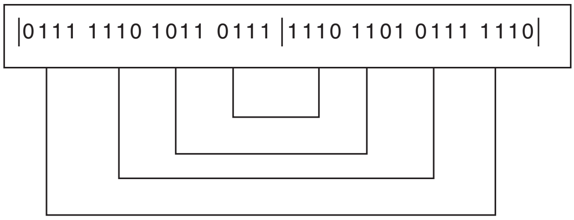

Fig. 17.4 An example of a rhythmic retrograde.

Highly experienced aficionados of serialist music could discriminate each of these from tone rows not exhibiting such symmetries—even when they believed they were literally guessing.

As Bharucha et al. indicate, melody is not the only aspect of music that can communicate structure. We have been investigating implicit learning of the rhythm of drum beats. In a training phase subjects listened to rhythms that instantiated a retrograde structure: The last half of each sequence was rhythmically the first half played backwards. Figure 17.4 shows an example: a '1' stands for a hit and a '0' for a rest. In a subsequent test phase using sequences constructed of novel chunks, when people said they were using intuition or just guessing, they reliably picked new retrogrades as being like the old ones: Figure 17.5 shows 95% confidence intervals for people's ability to discriminate the retrogrades from the non-retrogrades (d′). When trained subjects thought they were guessing or using intuition, they discriminated retrogrades from non-retrogrades (whereas untrained subjects did not).

Often the musical communication of structure involves feelings (for example, liking, intuition) that indicate the presence of structure that has been consciously discerned as existing, though not what that structure is (compare Norman, Price, & Duff, 2006). And sometimes people are not even aware they have discerned structure at all: this is

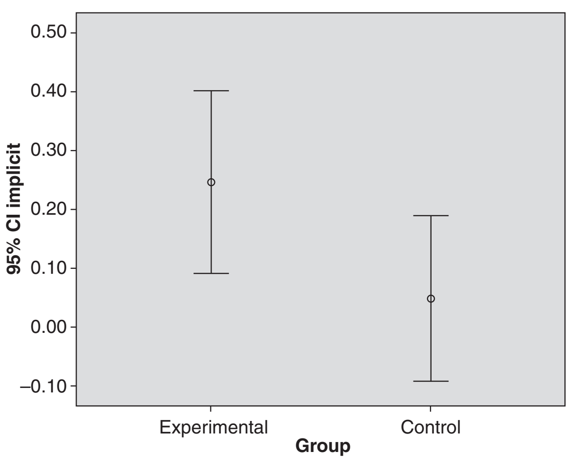

Fig. 17.5 Ability to discriminate rhythmic retrogrades from non-retrogrades when people thought they were guessing or just using intuition.

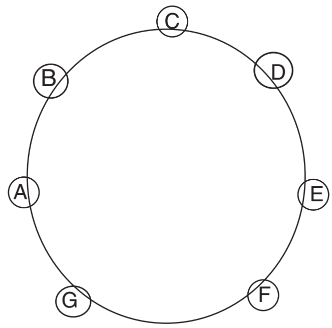

Fig. 17.6 Clock face structure of C major.

implicit learning, which underlies but does not yet fully constitute successful communication. Nonetheless, we had, as Bharucha et al. put it, induced an alignment of our subjects' brain states coding structure with those states we had intended.

Movement communicates affect and structure

Music and movement are similar. Melodies are movement in a pitch space (and also movement in a richer tonal space, e.g. Longuet-Higgins, 1976). Figure 17.6 shows the modulo representation of the tones of C major. A person physically moving around the clock face would produce movements structurally isomorphic with corresponding melodies. Like movement in tonal space, physical movements have natural opposites: moving two steps anti-clockwise is the opposite to moving two steps clockwise. A series of the opposites of a sequence of movements constitutes its inverse. We are currently exploring the implicit learning of symmetries in movement. We had people walk around a circle on the ground, walking out inverses (unbeknownst to them), under the cover story of practising moving meditation. This procedure is fairly time-consuming, so despite some initial encouraging results, we streamlined it by asking subjects to trace their fingers around a circle, as shown in Figure 17.7. Figure 17.7 shows two separate screen displays as seen by subjects. The subject first sees the left display and places their finger on the highlighted character. When the next display comes up (the one on the right), the subject moves their finger in the direction of the arrow to reach the next highlighted character. And so on until a sequence has been

Fig. 17.7 Two successive screen displays seen by people tracing out inversions.

traced out. None of our subjects knew Chinese characters; thus, the material was difficult to verbalize. We used a subset of the very same materials used in Kuhn and Dienes (2005). The melodies in the latter were all contained within the interval from middle C to the C above. To make the structure slightly harder to discern and the mapping onto music non-obvious, for our initial finger movement study, these two Cs were represented by different characters. To encourage processing the movements as a theme followed by a reply, after each half sequence had been traced, subjects tried to re-trace it unaided in one flow, then after the second half judged for the sequence as a whole whether it had occurred before. In the test phase, subjects classified the new sequences as rule following or not and then indicated the basis of their decision: guessing, intuition, rules or recollection (Dienes & Scott, 2005). As can be seen in Figure 17.8, when people thought they were guessing or using intuition, they picked out the inverses over the non-inverses, suggesting they had implicitly learned to detect inverses. Movement can communicate structure and hence feelings of intuitive rightness.

The methodology of our music and movement studies allows a novel test of the claim that 'music communicates movement'. Does listening to musical inverses induce a person to classify or like movement inverses over non-inverses? (Compare Altmann, Dienes & Goode, 1995, who used finite state grammars and obtained transfer between music and visual symbols.) To what extent is this true for other symmetries? Conversely, does following movement inverses prime one to like musical inverses? These remain questions for future research. They are related to the role of 'embodiment' in understanding language (e.g. Zwaan & Taylor, 2006). According to theories of embodied cognition, one understands the concept of an action or object by activating the motor patterns that are involved in engaging with that action or object. Is there in addition an embodied component to understanding syntax? We have shown that people can learn movement patterns instantiating grammars of complexity greater than context-free (namely, inverses and also retrograde inverses). Does understanding the inversion structure of music entail in any way activation of movement inversions? Is understanding centre embedding in language facilitated by performing movement retrogrades—or hearing musical retrogrades—or conversely harmed by performing movement inverses? These also remain for the time being unanswered questions.

Computational models of learning symmetries in music and movement

Learning inversions and retrogrades go beyond the computational models considered by Bharucha et al. The latter's models were designed to learn the co-occurrence relationships that occur in tonal music. However, inversions can occur in atonal music (as, for example, investigated by Dienes & Longuet-Higgins, 2004) and do not depend on any particular co-occurrence of tones or intervals. Similarly, learning inverses go beyond chunking models in the implicit learning field, because our test melodies were made of chunks not heard in the training phase (see Cleeremans & Dienes, 2008, for a review of computational models of implicit learning). What sort of model could learn the inversions we used? An inversion (or a retrograde) is a type of long distance dependency.

So any model that just learned to predict one tone (or interval) from the preceding one would fail. The model needs some kind of buffer or memory store. A computational model that has a long history of use in implicit learning is Elman's simple recurrent network, or SRN, illustrated in Figure 17.9 (see Cleeremans, 1993), which has an elegant memory store. This model learns structure in sequences not by having a memory of past events of any particular length but by (fallibly) learning how far into the past it should remember. At time one, the input units code the first note played and attempt to predict the second note at the output units by flowing activation through the hidden layer. Errors in prediction are used to adjust weights by back propagation. At the next time step, the hidden unit activations are copied to the context units. The input units code the second note and attempt to predict the third. But now the third is predicted by not just the input layer but also the activation across the context units, which carry information about the first note from the previous time step. So predictions about the third note can in principle be sensitive to the identity of both the first and second notes. Because the hidden layer receives input from the context and input units it now contains information about both the first and second notes. At the next time step this pattern of activation on the hidden layer is copied to the context units. The input units code the third note. Both the context units and input units together try to predict the fourth note. So the fourth note can in principle be predicted based on the first, second and third notes. And so on. Of course, whether the network really develops hidden unit representations carrying useful information from a long way back in time is in practice for any given case an open question. In principle it can do it if the small appropriate changes in the weights results in error reduction for each time step.

Cleeremans (1993) showed that initially when exposed to a sequence structured by a finite state grammar, the SRN learns first order dependencies—what elements predict what other immediate successors. It then learns to use the preceding two elements to predict the next—and the preceding three elements and so on. That is, the SRN learns chunks of progressively higher order. But there is no reason to think the SRN is just a chunk learner. If the buffer can keep track of number of tones it may eventually be able to learn the ith tone predicts the (i+ 4)th tone, regardless of the intervening material. Kuhn and Dienes (2008) trained the SRN on the materials of Kuhn and Dienes (2005) to see if it could indeed learn the long-distance dependencies that the inversion entailed.

We also used another model applied to implicit sequence learning by Cleeremans (1993), the fixed buffer model, illustrated in Figure 17.7. This model stores what has happened for a fixed amount of time into the past, here for precisely four notes into the past. This would correspond to say a fixed length working memory slave system (Baddeley, 1986). This model can learn associations between any tone in the preceding four trials and the current tone; in fact the model predicts any of these associations would be just as easy to learn as any other.

Kuhn and Dienes (2005) did not only present subjects with a test set consisting of novel chunks (the 'abstract' set) as we have discussed above; we also presented subjects with a test set where the tunes that happened to be inverses could be distinguished from the non-inverses by the frequency with which the chunks had occurred in training. Thus we can see whether subjects or models are more sensitive to the inversion or to chunk statistics.

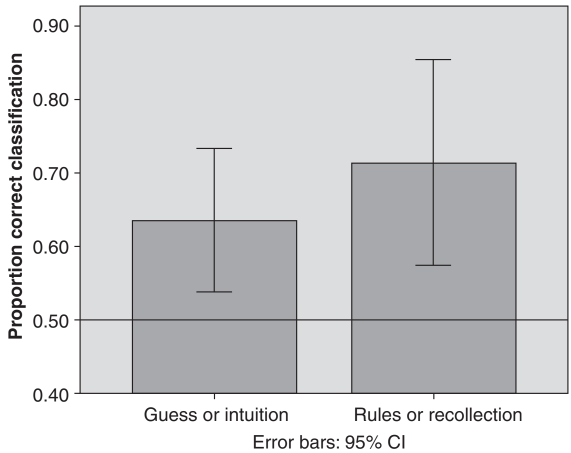

Fig. 17.8 Ability to discriminate movement inverses from non-inverses when people thought they were guessing or just using intuition.

We trained the models on the same material as people for the same number of exposures and tested them on the same test sets. Figure 17.10 shows the ability of each model to discriminate inverses from non-inverses where each point corresponds to a certain combination of parameter values. There were a number of parameters free to vary as is typical for any computational model: the learning rate, the number of hidden units, and so on. Because different combinations of free parameter values can dramatically affect the results the model produces, how do we know when a model is a good explanation? If a model could predict all possible outcomes by adjusting free parameters it would hardly be an explanation of any outcome. To determine the explanatory power of each model we explored a full range of parameter space to determine to what extent model predictions concentrated in the region of human performance. We regard this as a superior methodology for evaluating models than simply fitting data.

The y-axis in Figure 17.10 shows performance on the test set where discrimination could be achieved by knowledge of chunks and the x-axis shows performance on the abstract set where all chunks were novel. The small squares are specific combinations of free parameter values for fixed buffer models and the crosses are for SRN models. It can be seen that the SRN models were very sensitive to chunks; the crosses rise steeply

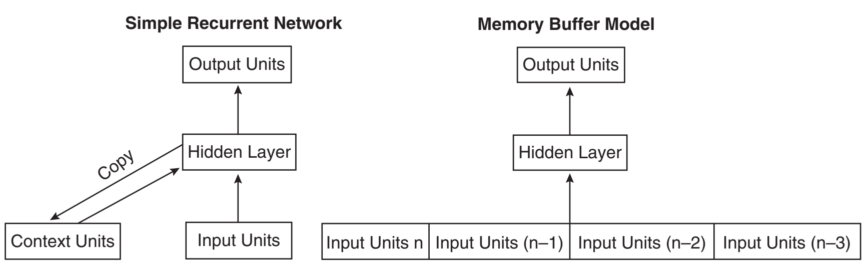

Fig. 17.9 The two computational models used by Kuhn and Dienes (2008).

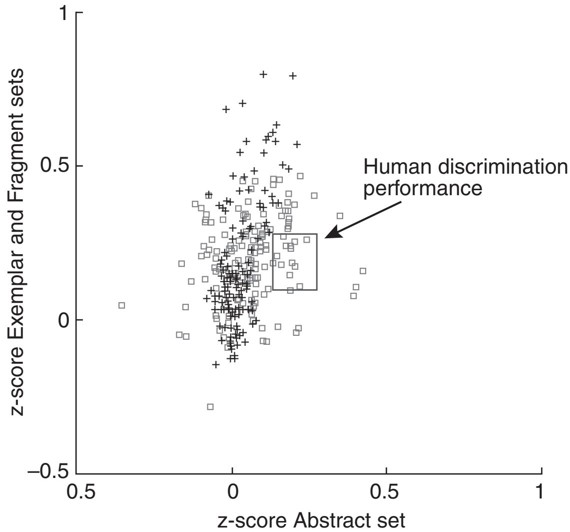

Fig. 17.10 Discrimination performance of models and people. The y-axis represents ability to discriminate inverses from non-inverses when chunk frequency can aid that discrimination. The x-axis represents discrimination performance when all chunks in the test phase were novel. Small squares represent fixed buffer models; crosses SRN models. The large square is the mean for people plus or minus one standard error.

in the figure. Conversely, the fixed buffer models are not sensitive to chunks in particular: As long as there are associations, they don't have to be between adjacent elements to be learned. The large square is the mean for human data plus or minus a standard error. Notice that while the SRN *can* learn the abstract set and the chunk sets as well as people for particular combinations of free parameters, what the SRN typically does is learn chunks. In consequence, while 12 out of 150 fixed buffer models are contained in the large square, zero out of 150 SRN models are. More fixed buffer models are concentrated around human data than SRN models are ($p < 0.0005$); characteristically, the buffer model behaves more like people than the SRN model does.

The models learned particular long distance associations. But this is different from learning inverses as such. To know that a C predicts an F four tones later is not to detect a mirror symmetry as such. We do not yet know what people have implicitly learned either—a set of specific associations or the 'operation over variables' (Marcus, 2001) that whatever interval was in ith position in the theme must be (−1) times that interval in the ith position in the inverse. The latter knowledge allows generalizations to C predicting tones other than specifically F where appropriate, and to spotting inverses of arbitrary length. Our models are models of learning the specific task we set subjects but not yet models of learning the full symmetry our materials instantiated. We look forward to future work determining what it is that people actually learn and what models can explain that ability. Finally, our models fail completely to be models of musical interpretation in general. What we need is a model that explains the tonal nature of our musical appreciation, as addressed by the model of Bharucha et al., as well as our ability to detect and appreciate symmetry in music, as we have taken first step to addressing with the models reported here.

References

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

Baddeley, A.D. (1986). *Working Memory*. Oxford: Clarendon Press.

Brooks, L. R. (1978). Non-analytic concept formation and memory for instances. In E. Rosch & B. Lloyd (Eds.), *Cognition and concepts*. Hillsdale, NJ: Erlbaum.

Cleeremans, A. (1993). *Connectionist models of sequence learning*. Cambridge, MA: MIT Press.

Cleeremans, A., & Dienes, Z. (2008). Computational models of implicit learning. In R. Sun (Ed.), *Cambridge Handbook of Computational Psychology*. Cambridge: Cambridge University Press

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

Dienes, Z. (2008) Subjective measures of unconscious knowledge. *Progress in Brain Research*, *168*, 49–64.

Dienes, Z., & Kuhn, G. (submitted). Implicitly learning to detect symmetries: Reply to Desmet et al. http://www.lifesci.sussex.ac.uk/home/Zoltan_Dienes/Dienes & Kuhn 2008.doc

Dienes, Z. & Longuet-Higgins, H. C. (2004). Can musical transformations be implicitly learned? *Cognitive Science*, *28*, 531–58.

Dienes, Z., & Scott, R. (2005). Measuring unconscious knowledge: Distinguishing structural knowledge and judgment knowledge. *Psychological Research*, *69*, 338–51.

Jamieson, R. K. & Mewhort, D. J. K. (2009). Applying an exemplar model to the artificial-grammar task: Inferring grammaticality from similarity. *Quarterly Journal of Experimental Psychology*, *62*, 550–75.

Kuhn, G., & Dienes, Z. (2005). Implicit learning of non-local musical rules. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *31*, 1417–32.

Kuhn, G., & Dienes, Z. (2008). Learning non-local dependencies. *Cognition*, *106*, 184–206.

Longuet-Higgins, H. C. (1976). The perception of melodies. *Nature*, *263*, 646–53.

Norman, E., Price, M.C., & Duff, S.C. (2006) Fringe consciousness in sequence learning: The influence of individual differences. *Consciousness and Cognition*, *15*, 723–60.

Marcus, G. (2001). *The algebraic mind*. Cambridge, MA: MIT Press.

Mozer, M. C. (1994). Neural network music composition by prediction: Exploring the benefits of psychoacoustic constraints and multi-scale processing. *Connection Science*, *6*, 247–80.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, *119*(3), 264–75.

Pothos, E. M. (2007). Theories of artificial grammar learning. *Psychological Bulletin*, *133*(2), 227–44.

Pothos, E. M. & Bailey, T. M. (2000). The importance of similarity in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 847–62.

Reber, A.S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*, 317–27.

Servan Schreiber, E., & Anderson, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(4), 592–608.

Seth, A., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L. (2008). Measuring consciousness: relating behavioural and neurophysiological approaches. *Trends in Cognitive Sciences*, *12*, 314–21.

Shanks, D. R. (2005). Implicit learning. In K. Lamberts & R. Goldstone (Eds.), *Handbook of cognition* (pp. 202–20). London: Sage.

Zizak, D. M. & Reber, A. S. (2004). The structural mere exposure effect: The dual role of familiarity. *Consciousness and Cognition*, *13*, 336–62.

Zwaan, R.A., & Taylor, L.J. (2006). Seeing, acting, understanding: motor resonance in language comprehension. *Journal of Experimental Psychology: General*, *135*, 1–11.