Perceiving and Learning Harmonic Structure: Some News from MUSACT

Barbara Tillmann * & Jamshed J. Bharucha **

* Université de Bourgogne LEAD-CNRS
Boulevard Gabriel
F-21000 Dijon, France
tillmann@u-bourgogne.fr
** Dartmouth College, Department of Psychology
Hanover, NH 03755, USA
Jamshed.J.Bharucha@dartmouth.edu

Abstract

Harmonic priming research provides evidence that local and global harmonic contexts influence processing of target chords for musician and nonmusician listeners. Perceptual analysis of musical structures partly depends on how listeners' knowledge of tonal hierarchies are represented in the mind. Internalized representations of structural regularities generates musical expectations and facilitates the processing of harmonically related events. MUSACT (Bharucha, 1987) provides a connectionist framework for the representation of tonal knowledge. Activation spreading through a network of representational units accounts for the influence of local and global context on the processing of chords. The listeners' implicit knowledge of harmonic structures is acquired through mere exposure to the conventional relationships between musical events in Western tonal music. MUSACT represents the idealized end-state of such a learning process. In this paper, we present computer simulation demonstrating that MUSACT's basic structure can be learned by mere exposure via self-organization.

Keywords: harmonic priming, global and local context effects, spreading activation model, unsupervised learning of harmonic structure

1 Perceiving Harmonic Structure

1.1. MUSACT: A Connectionist Model of Western Tonal Knowledge

Several experimental studies have shown that listeners have internalized the tonal-harmonic hierarchy through passive exposure to Western tonal music (Bharucha, 1987; Francès, 1958; Krumhansl, 1990). This internalized knowledge is activated by a musical context and generates expectancies for related events to follow. The MUSACT model (Bharucha, 1987) provides a connectionist framework for understanding how this knowledge is represented and how expectations are generated. In this model, knowledge of Western harmony is conceived of as a network of interconnected units which are organized in three layers: tones, chords, and keys. There are twelve tone units (representing the twelve pitch classes), twelve major chord units, twelve minor chord units, and twelve major key units. Each tone unit is connected to the chord units representing chords of which that tone is a component. Analogously, each chord unit is connected to the three major key units representing keys of which it is a member. The structure of the Western musical system is expressed in the model by the strength of the connections that link tone units to chord units and chord units to key units.

When three triadic tones are played (say c-e-g), the units representing these tones are activated, and phasic activation spreads to the chord units via the connected links. (The phasic activation of a unit is its change of activation from the previous iteration to the present one.) The chord unit connected to all three tones receives the strongest activation (the C major chord in this example). During a second cycle, phasic activation from the active chord units spreads towards the key units (bottom-up activation) and back down to the tone units (top-down activation). During the next cycle, activated key units send top-down activation to chord units that simultaneously received bottom-up activation from the tone units. After several cycles, the model reaches a state of equilibrium. In the equilibrium state, the activation pattern reflects the Western tonal hierarchy at each level of the network. In this example, the C major chord unit has the highest activation, followed by the F major and G major chord units. Activation decreases with increasing distance around the cycle of fifths (Figure 1). The activation pattern of chord units represents the degree to which each chord is expected, and accounts for the facilitation of the processing of related chords. The more a chord unit is activated, the more that chord is expected and its processing facilitated if it occurs.

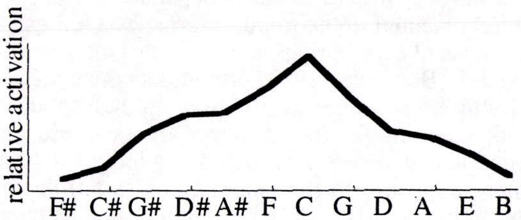

Fig. 1: Relative activations observed for major chord units once MUSACT has reached equilibrium after the presentation of a C major chord

MUSACT also addresses the building up of harmonic expectancies over time. Once the model has reached equilibrium after an event, the pattern of activation decays over time. If another event is presented to the system, its activation is added to the residual activation from the previous event. The activation of a unit i in the network is a function of not just the most recent event e, but also of the previous event, e-1, the activation of e-1 being itself a function of event e-2 , and so on. The total activation, $a_{i,e}$, of a unit i is given by the following equation:

$$a_{i,e} = a_{i,e-1}\,(1-d)^t + A + \sum_{c=1}^{q} \Delta a_{i,e,c} \qquad (1)$$

where d represents the rate (varying between 0 and 1) at which activation decays following the offset of the last event, t represents the time transpired since the last offset, A the stimulus activation, and the third term the total phasic activation of unit i in response to event e, accumulated over q reverbatory cycles that are necessary to reach equilibrium. The total activation, $a_{i,e}$, of a unit i (a tone, a chord or a key) after an event e is then an additive function of three quantities: (1) the decayed activation caused by the previous event e-1, (2) the bottom-up activation caused directly by the stimulus itself (i.e. the tones), and (3) the indirect activation received from other units in response to event e. The first quantity represents the global context, the second quantity represents the stimulus effect itself, and the third the local context (the most recent event). The activations due to several chords are accumulated as the sequence unfolds, yielding an aggregate expectation for further incoming events. In this way, MUSACT takes into account the development of expectations in long harmonic contexts.

1.2. MUSACT: A Spreading Activation Account of Harmonic Priming

Harmonic priming research explores the influence of a previous harmonic context on the processing of upcoming events. A harmonic context generates expectancies and primes chords that are harmonically related to the context. MUSACT explains the development of these expectations via activation spreading through a network representing tonal knowledge. After the presentation of a harmonic context, the activation pattern of chord units reflects the expectancy for following events. The context can be a single chord or entire chord sequences whose activations are accumulated over time. The extent to which a target chord is primed is a function of the activation of the corresponding chord unit. Simulations of MUSACT can be compared with priming data from human subjects.

1.2.1. Single Chord Priming

In these studies (Bharucha & Stoeckig, 1986, 1987; Tekman & Bharucha, 1982), participants heard a prime chord followed by a target chord. The prime and target were either closely related (belong to the same key) or distantly related harmonically. For example, if the prime chord was C major, Bb major would be a related target and F# major an unrelated target. On half of the trials, the target chord was slightly mistuned, and participants were asked to make a speeded intonation judgment, i.e., to decide as quickly as possible whether the target chord was in tune. The priming effect was shown by (1) a bias to judge targets to be in tune when they were related to the prime, and (2) shorter response times for in-tune targets when they were related to the prime, and for out-of-tune target when they were unrelated to it. Thus, a single chord can generate expectancies for related chords to follow, resulting in greater consonance and faster processing for expected chords. This outcome provides support for the model. The activation pattern of chord units simulates harmonic expectations of human subjects and accounts for the facilitation of the processing of related chords. The more a chord unit is activated, the more the chord is expected and its processing facilitated.

In the priming studies cited above, the related prime-target pairs shared more tones in common than did the unrelated target-pairs, leaving open the possibility that priming is driven solely by common features and doesn't require the top-down influences based on prior knowledge. Tekman and Bharucha (1998) tested this possibility by pitting shared tones against conventional relatedness. Two types of target were selected: one was more psychoaccoustically similar to the prime, the other more closely related on the basis of harmonic convention. For example, a C major prime shares a tone with an E major target but does not share a tone with a D major target; yet D major is more closely related to the prime in conventional usage. In the model, the pattern of activation changes qualitatively during the reverberatory cycles from initial activation to equilibrium. In early activation cycles (bottom-up activation), the pattern of activation reflects the number of tones shared by the prime and the target, whereas at equilibrium (after top-down influences have had their effect) the pattern reflects conventional relatedness - distance around the cycle of fifths. Results revealed facilitation for psychoacoustically similar targets when they followed after a short (50 ms) stimulus onset asynchrony (SOA), and facilitation for conventionally related targets after a longer SOA (500ms or longer). While both psychoacoustic similarity and conventional relatedness drive priming, the influence of the former is short-lived, precisely as predicted by temporal course of activation in the model.

1.2.2. Priming in Long Sequences of Chords

In MUSACT, the activations due to a sequence of chords are accumulated as the sequence unfolds. The total activation pattern reflects the harmonic hierarchy of the sequence's underlying key. The position of the target chord in this hierarchy (i. e., its activation level) determines the degree of expectation for this event. Recent studies extended harmonic priming effects to large contexts. Bigand and Pineau (1997) manipulated the global context of eight-chord sequences. Expectations for the last chord (the target) were varied by changing the harmonic context created by the first six chords. The last two chords were held constant. In the expected condition, the last chord was a harmonically stable tonic chord, part of an authentic cadence (V-I). In the unexpected condition, the last chord took the form of a less stable fourth harmonic degree following an authentic cadence (I-IV). Participants were faster and more accurate in their intonation judgment of the last chord when it was expected. These results suggest that harmonic priming involves higher level harmonic structures and does not occur only from chord to chord. In particular, priming reflects global harmonic contexts as well as just the local effect of the previous chord.

In simulations of the model performed with the sequences of Bigand and Pineau (1997), the first seven chords defined the prime. The resultant activation patterns for major chords were interpreted as the array of expectations for the major chords to follow (Figure 2). The target chord unit received stronger activation when it acted as a stable tonic chord (I) in the expected context than when it was a less stable sub-dominant chord (IV) in the unexpected context. The activation pattern in the neural net thus takes into account the influence of the global context and mimics human performance, showing a facilitation of the expected targets in comparison to the unexpected targets.

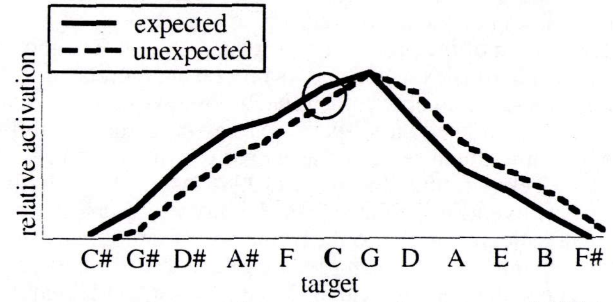

Fig. 2: Relative activations observed for major chord units once MUSACT has reached equilibrium on the penultimate chord in the expected and unexpected contexts. For convenience, the state of the network is represented with reference to the C major key. In this key, the C chord unit represents the target.

Global harmonic priming was extended to wider harmonic contexts in a recent study (Bigand, Madurell, Tillmann & Pineau, in press). The global context was manipulated in 14 chord sequences at three levels, while holding constant the chord prior to the target (local context). The function of the target chord was changed by transposition. In the highly expected condition, the whole sequence is played in the same key, and the target chord is part of an authentic cadence (V-I) that closes the overall structure. In the unexpected condition, it is played in the dominant key and the target chord is the fourth harmonic degree following an authentic cadence (I-IV). These two conditions replicated those of Bigand and Pineau (1997) with longer chord sequences. In the middle expected condition, the first half is harmonically identical to the first half of the highly expected condition and the second half to that of the unexpected condition. Although the chords of the second half are strictly identical, the target chord in the middle expected condition is no longer the fourth harmonic degree following an authentic cadence. In this context it may be analyzed as part of the authentic cadence (V-I) that returns to the main key. The results provide evidence that musical expectations derive from various levels of hierarchical structure. Strongest facilitation was observed for highly expected targets, as the target chord was expected at both high and intermediate levels. Facilitation was reduced when it was expected at the higher level only (i.e. middle expected condition). The weakest priming effect was observed when the target chord was not strongly expected at both high and intermediate levels.

Simulations run with the long chord sequences are globally in accordance with human performance. The target chord was less activated in the unexpected context than in the middle expected context, and in the middle expected context it was less activated than in the highly expected condition. MUSACT thus accounts for subtle effects of large scale structures in music. The neural net keeps some trace of the first key until the end of the sequence because activations present at the beginning of the sequence are added to those created by the new key.

1.2.3. Local and Global Harmonic Priming

To what extent does interposing a harmonically unrelated chord between the global context and the target weaken the priming effect? To what extent does inserting a harmonically related chord before the target compensate for the lack of any global harmonic relation between the target and the sequence? Tillmann, Bigand & Pineau (in press) varied the target's relatedness on a global and local level and performed crude changes in harmonic relationships at both global and local levels. For example, in a C major key, the target chord was globally and locally related (GRLR) when it was a tonic chord (C) and was preceded by a dominant chord (G). It was globally related but locally unrelated (GRLU) when the preceding dominant chord was played one semitone higher (G#). In this case, the target and the preceding chord do not belong to the same key. The target was globally unrelated but locally related (GULR) when only the first six chords of the sequences were transposed one semitone above (i. e. in the C# major key). Here the key of the first six chords is weakly related to the keys of the target chord and its preceding chord (i.e., C and G major keys). Finally, the target chord was both globally and locally unrelated (GULU) when the first seven chords were transposed one semitone above (in the C# major key).

For this experimental material, simulations with MUSACT predict that both local and global context influence harmonic priming, with the strength of the global context depending on the tempo. During the ongoing sequence, activations of both local and global context add, and decay (exponentially) over time, with the most recent events being the most active. The strength of this decay varies as a function of the tempo. As a consequence, the effect of global context should be more pronounced at a fast tempo. Simulations were conducted with the first seven chords of each sequence for two tempi. The activation of the target chord unit depended on whether one or two sources of priming are present. It was the highest for the GRLR condition, as both contexts were related to the target chord. Activation decreased for GRLU and GULR respectively, with only one context related to the target. It was the lowest for GULU, where the target chord had no relation to the previous context. At a slow tempo, the global and the local contexts exerted roughly similar effects. At a fast tempo, however, the global context strongly prevailed over the local context. The performance of participants thus demonstrated a strong effect of both global and local context. Target chords were processed more accurately and quickly when they were locally or globally related to the previous context. In accordance with MUSACT's predictions, the effect of global context tended to be more pronounced at a fast tempo. In sum, the simple accumulation of tonal hierarchy patterns takes into account the priming effects observed for local and global contexts. The influence of tempo also suggests that tonal hierarchy patterns are added and weighted by decay. Priming effects seems to be the result of activation spreading via a stable cognitive structure that links related chords.

1.3. MUSACT: End-state of Learning

All harmonic context effects summarized above were observed independently of the extent of musical expertise. Data from both musicians and nonmusicians fit with the predictions of the spreading activation model. Harmonic priming thus seems to reflect an underlying system that can be acquired without formal instruction, presumably through passive exposure to Western music, in which constraints on harmonic relationships are pervasive. MUSACT represents an idealized end-state of such a learning process, as it is based on these music theoretic constraints (Bharucha & Olney, 1989). Its connectionist representation of tonal knowledge is a powerful framework for understanding the influence of context on harmonic expectations. However, it is a constraint-satisfaction model, and a crucial point is to analyze how such a representation of tonal knowledge is learned by mere exposure to musical material. It has been suggested that MUSACT can be learned by unsupervised learning. Unsupervised learning mechanisms extract underlying regularities of the tonal system, i. e., co-occurrence of notes in chords or of sets of chords in keys (Bharucha, 1991, 1992). In the following, computer simulations are presented that were run with Self-Organizing Maps (Kohonen, 1995), an unsupervised learning algorithm.

2 Learning of Harmonic Structure by Self-Organization

2.1. General Principles of Self-Organization

The Self-Organizing Map algorithm is an unsupervised learning algorithm that creates topological mappings between the input data and map units. For two similar input patterns, the responding map units are located near each other on the map. This algorithm is based on principles of cortical information processing, in particular the formation of spatial ordering in sensory processing areas (i.e. somatosensory, vision and audition). In the primary visual cortex, the orientation of stimuli to which cells respond best changes in an orderly fashion across the context: nearby cells respond best to similar orientations (Hubel & Wiesel, 1962). The auditory cortex displays a topographical organization in which cells responding best to different frequencies are arranged in orderly fashion (Brugge & Reale, 1985; Wessinger, Buonocore, Kussmaul & Mangun, 1997).

The SOM algorithm has been applied in a wide variety of fields: starting from neurophysiological research over physics, signal and data processing to speech analysis and recognition. It has also been used to model various aspects of music (e. g.

Gjerdingen, 1990; Griffith, 1994, 1995; Leman, 1995). Our principal aim here was to simulate the learning of a representation of harmonic knowledge with MUSACT's properties.

In SOMs, the training is done in an unsupervised manner as the network changes its connections based on the properties of the input data. Unsupervised learning algorithms seem to be the closest to real music perception as no external teacher gives feed-back on the organization of chords or tonalities. Competitive learning is one algorithm for data-driven self-organized learning. It represents a process in which neural network units gradually become sensitive to different input stimuli or categories (Rumelhart & Zipser, 1985). The specialization takes place by competition among the units: when an input arrives, the unit that is best able to represent it, wins the competition and is allowed to learn the representation even better (as will be described below). If there exists an ordering between the units, i.e. the units are located on a discrete lattice (the self-organizing map SOM) the competitive learning algorithm can be generalized: not only the winning unit, but also its neighbors are allowed to learn. Neighboring units will gradually specialize to represent similar inputs and the representation becomes ordered on the map. SOM produces a mapping from an n-dimensional data space onto a two-dimensional space, which is represented by a grid of neuron-like units. After learning, each unit is specialized to detect a particular input pattern, and a topological organization of the input data can be discovered on the grid, such that similar input patterns activate nearby units.

The network that develops a self-organized map consists of two layers of units: an input layer and a two-dimensional grid layer. These two layers are fully interconnected by synapses, i. e. each grid unit has a synapse feeding into it from each input unit. Prior to learning, the connection strengths (the weights) are initialized to random values. When a stimulus is presented, the input units, *i*, tuned to its features are activated. These activations spread via the connected links, *w(i,j)*, to the grid layer. Each unit, *j*, of the second layer accumulates the activation it receives from the input units. The activation of each grid unit *j* is given by:

$$a(j)=\sum_{i} a(i) * w(i,j) \qquad (2)$$

The unit *j* with the highest activation (i. e., the winning unit) is selected. During the learning phase, the associated weight vectors of the winning unit and those within a neighborhood set *N* are updated. The weights of units outside the neighborhood set are kept constant. Learning consists of updating the weights feeding into the winning unit and its neighbors with the following algorithm (Kohonen, 1995) :

$$w(t+1)=\frac{w(t) + \eta(t) * a(t)}{\|w(t) + \eta(t) * a(t)\|} \qquad (3)$$

where $w(t+1)$ is the weight vector at time $t+1$, $w(t)$ at time t , and $\eta(t)$ is the learning rate. This learning rule moves the weight vectors closer to the input vector, making the winning unit and its neighbors more likely to win the competition when presented with this input or ones similar to it.

The neighborhood set *N* is set to be wide at the beginning of learning. During learning, it decreases monotonically until it consists of the winning unit alone. As learning begins, a large neighborhood allows a global organization to take place. With a smaller neighborhood radius, the units become adapted to the individual patterns and its close relatives, and a local organization takes place.

Before learning, there is no particular organization among the grid units. When the net is trained by repeated presentation of the input data, it begins to self-organize.

Topographic pattern begins to appear, such that units that are topographically close in the array will be activated by similar input stimuli. SOM can be conceived of with one grid layer or be adapted to multilayer hierarchical self-organizing maps (HSOM) (Lampinen & Oja, 1992).

2.2. Learning of Harmonic Structure

Our principal aim is to analyze how MUSACT's basic structure can be learned by mere exposure via self-organization. A three layer hierarchical system is defined: the input layer consists of 12 units, the second layer is a map of 36 units and the third layer is a map of 16 units. The input units are tuned to the 12 chromatic scale tone units that represent octave-equivalent pitch categories. The second layer will learn to specialize in the detection of chords and the third layer in the detection of groups of chords defining a key. In the input layer, a more abstract coding than just frequency is chosen as it has been shown that neural net models can learn octave equivalent pitch classes (Bharucha & Mencl, 1996). The input unit is activated if the corresponding tone to which it is tuned occurs in the chord, and 0 otherwise. The units of the first and second layers are fully interconnected via a connection matrix; and the units of the second and third layers with a second connection matrix. Before learning, the strengths of all connections are initialized to random values.

Several constraints are applied to the learning material to which the system is exposed. These constraints reflect the restrictions inherent in MUSACT, and should favor the learning of its structure. If this structure can be learned, the learning material can be extended to more complex patterns in further simulations.

In the present simulations, the learning set is restricted to 24 chords (12 major and 12 minor chords) and 12 major keys. An important constraint is the composition of major keys: a major key is defined by a group of six chords (three minor and three major chords) presented to the input layer one by one without decay.

The training patterns are presented in random order during each training cycle. Learning consists of two phases. In the first phase, the second layer is trained by the presentation of 24 chords (12 major chords, 12 minor chords) presented individually. In the second learning phase, the third layer is trained with sets of six chords representing major keys. A major key is defined by three minor and three major chords, e.g., the C major key is represented by the major chords C, F, and G, and the minor chords d, e, and a. These six chords are presented individually to the input layer. For each input chord, the best matching unit is chosen from the second-layer map and its index b is stored in memory until the end of the presentation of the chord set. The pattern of indexes b (without decay) defines the input for the training of the third layer.

At the beginning of learning, the neighborhood radius is set to its maximum and decreases during training until it reaches 0, i. e. only the winner learns. The learning rate is kept constant at the beginning of learning. In the convergence phase (i.e. when only the winning unit learns), the learning rate decreases over the training cycles.

For both training sessions, the weight changes decrease over the training cycles and as the neighborhood decreases. When weights have converged to practically stationary values, the maps are calibrated in order to locate images of inputs on them. During learning, units are specialized for the detection of chords and for the detection of sets of chords (referred to henceforth as 'keys').

The calibration phase reveals a topographic organization of both maps. Chord units in the second layer are organized so that neighboring units share component tones. Chords that do not share tones are segregated; for example, D# major, d# minor, B major, b minor, G major and g minor are on one side of the map while F major, f minor, A major, a minor, C# major, and c# minor are on the other side. In the third layer, keys sharing chords and tones are represented by units close to each other on the map. The organization of the specialized key units represents the cycle of fifths, a music theoretic concept. Musical distances between keys are represented on a circle with keys sharing all but one notes as neighbors. This organization emerges in the third layer after learning with groups of chords.

The weights of both connection matrices have changed after learning. When considering the links feeding into winning units, the learned matrices reflect the predefined links of the MUSACT model. Each tone unit of the input layer has six connections to the winning units of the chord layer, i. e. to the six chords of which it is a part. Each chord unit is linked to three key units in the third layer, i.e. to the keys to which it belongs. The connections defined by music theory in MUSACT are thus learned by self-organization.

The present simulation provided evidence that a representation of tonal knowledge can be learned by self-organization. Without external feedback or supervision, the structure of the material to which the system is exposed to is learned in the connection matrices. As a consequence of these changed connections, units specialize in the detection of chords and keys. Interestingly, both maps reveal a topographic organization. Units responding to similar stimuli (i.e. chords or groups of chords) are located in neighborhood on the map.

Further simulations will be necessary to test the learned model as both a feedforward and a reverberation system and to compare its behavior to that of MUSACT. When the new network is used as a simple feedforward system, tone units send their activation to the second and third layers, and the activation levels of these two layers represent the output. This feedforward activation incorporates psychoacoustic information only (tones present in the stimuli) without any top-down influences. Adding reverberation permits a top-down influence of the key layer on the chord and tone layers. In the MUSACT model, reverberation changes qualitatively the pattern of activation because of this top-down influence of learned, schematic structures. Further extensions of simulations will be the use of more ecologically valid materials, as real chord sequences or a richer coding based on subharmonics (Parncutt, 1988, 1989).

3 Conclusion

Harmonic priming studies provide evidence that Western listeners have internalized an implicit knowledge of regularities of the Western tonal system. This knowledge is activated by a context and creates expectations for subsequent events. The processing of expected events is facilitated in contrast to unexpected events. Self-organizing algorithms can account for the passive perceptual learning of tonal knowledge by mere exposure. During learning, specialized representational units are formed for combinations of musical events (tones, chords) that occur with great regularity. Results reveal that the MUSACT model, constructed on the basis of music theoretic constraints and originally proposed as an idealized end-state of a learning process, can be learned by self-organization.

References

Bharucha, J. J. (1987). Music cognition and perceptual facilitation: a connectionist framework. Music Perception, 5, 1-30.

Bharucha, J. J. (1991). Pitch, harmony and neural nets: A psychological approach. In P. Todd & G. Loy (Eds.), Music and Connectionism. Cambridge: MIT Press, pp. 84-99.

Bharucha, J. J. (1992). Tonality and learnability. In M. R. Jones & S. Holleran (Eds.), Cognitive bases of musical communication, Washington, DC: American Psychological Association, pp. 213-223.

Bharucha, J. J. & Mencl, W. E. (1996). Two Issues in auditory cognition: Self-organization of octave categories and pitch-invariant pattern recognition. Psychological Science, 7, 142-149.

Bharucha, J. J. & Olney, K. L. (1989). Tonal cognition, artificial intelligence and neural nets, Contemporary Music Review, 4, 341-356.

Bharucha, J. J., & Stoeckig, K. (1986). Reaction time and musical expectancy: Priming of chords. Journal of Experimental Psychology: Human Perception & Performance, 12, 403-410.

Bharucha, J. J. & Stoeckig, K. (1987). Priming of chords: Spreading activation or overlapping frequency spectra? Perception and Psychophysics, 41, 519-24.

Bigand, E., Madurell, F., Tillmann, B., & Pineau, M. (in press). Effect of global structure and temporal organization on chord processing. Journal of Experimental Psychology: Human Perception & Performance.

Bigand, E. & Pineau, M. (1997). Global context effects on musical expectancy. Perception & Psychophysics, 59, 1098-1107.

Brugge, J. F. & Reale, R. A. (1985). Auditory cortex. In A. Peters & E. G. Hones (Eds.), Cerebral Cortex: Association and Auditory Cortices. New York: Plenum Press.

Francès, R. (1958). La perception de la musique. Paris: Vrin, Transl. W. J. Dowling, (1988), The perception of music, Hillsdale, NJ: Erlbaum.

Gjerdingen, R. O. (1990). Categorization of musical patterns by self-organizing neuronlike networks. Music Perception, 8, 339-370.

Griffith, N. (1994). Modeling the acquisition and representation of musical tonality as a function of pitch-use through self-organizing neuronlike networks. Unpublished doctoral thesis, University of Exeter, Dept. of Computer Science.

Griffith, N. (1995). Connectionist visualization of tonal structure. AI Review, 8, 393-408.

Hubel, D. H. & Wiesel, T. N. (1962). Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. Journal of Physiology, 160, 106-156.

Kohonen, T. (1995). Self-Organizing Maps. Springer: Berlin.

Krumhansl, C. L. (1990). Cognitive foundations of musical pitch. Oxford University Press.

Lampinen, J. & Oja, E. (1992). Clustering properties of hierarchical self-organizing maps. Journal of Mathematical Imaging and Vision, 2, 261-272.

Leman, M. (1995). Music and Schema Theory. Springer: Berlin.

Parncutt, R. (1988). Revision of Terhardt's psychoacoustical model of the root(s) of a musical chord. Music Perception, 6, 65-94.

Parncutt, R. (1989). Harmony: A psychoacoustical approach. Springer: Berlin.

Rumelhart, D. F. & Zipser, D. (1985). Feature discovery by competitive learning. Cognitive Science, 9, 75-112.

Tekman, H. G. & Bharucha, J. J. (1992). Time course of chord priming. Perception & Psychophysics, 51, 33-39.

Tekman, H. G. & Bharucha, J. J. (1998). Implicit knowledge versus psychoacoustic similarity in priming of chords, Journal of Experimental Psychology: Human Perception and Performance, 24, 252-260.

Tillmann, B., Bigand, E., & Pineau, M. (in press). Effects of local and global contexts on harmonic expectancy. Music Perception.

Wessinger, C. M., Buonocore, M. H., Kussmaul, C. L., & Mangun, G. R. (1997). Tonotopy in human auditory cortex examined with functional magnetic resonance imaging. Human Brain Mapping, 5, 18-25.

NOTE

This research was supported by a grant from the National Science Foundation (SBR-9601287) to the second author.