9th International Conference on Music Perception and Cognition

Alma Mater Studiorum University of Bologna, August 22-26 2006

An interval cycle-based model of pitch attraction

Matthew Woolhouse
Centre for Music & Science,
Faculty of Music,
Cambridge University, UK
mhw31@cam.ac.uk
mwoolhouse@mac.com

Ian Cross
Centre for Music & Science,
Faculty of Music,
Cambridge University, UK
ic108@cam.ac.uk

ABSTRACT

A model of pitch attraction based upon interval cycles is presented. The components of the model are formally described and illustrated using a simple two-chord example (G to C). Two experiments are used to test the model. In the first experiment, *Chord to chord*, subjects rated the level of attraction/resolution between pairs of chords. In second experiment, *Chromatic chord to probe tone*, subjects rated the level of attraction from chords to probe tones. Both experiments produced results that significantly agreed with the model's predictions of pitch attraction for temporally contiguous musical events.

Keywords

Pitch attraction, interval cycles, harmony, probe tone.

INTRODUCTION

The notion that relationships between musical pitches can be conceived of, or experienced, in terms of forces has been widespread in the musical and musicological literature since at least the work of Kurth (1931), particularly, for example, the idea that one pitch may be 'attracted' towards another. While the idea is generally treated metaphorically within music theory, it has been a focus of theoretical and empirical research within the cognitive sciences of music, typically being approached in terms of the modulation of tension in the ongoing experience of music (Lerdahl & Jackendoff, 1983; Krumhansl, 1996; Larson, 1998; Lerdahl, 2001).

Many different frameworks have been proposed to account for the tendency of certain pitches (such as the leading-note or the subdominant) to 'move towards', 'be attracted to' or 'resolve onto' certain other pitches (the leading-note tending to resolve to the tonic, the subdominant resolving to the mediant). The principles of Auditory Scene Analysis (ASA) elucidated by Bregman (1991) do seem to account well for much voice-leading practice in western music (see Huron, 2001), while the associationist approach exemplified in the works of Tillmann (2000), Krumhansl (2000) and Oram (1995) propose that the basis of such tendencies is statistical (frequently-occurring pitches being likely to be 'resolved to'). Psychoacoustical approaches (such as those of Terhardt, 1974, and of Parncutt, 1989) suggest that events that are perceptually more stable (such as root position triads) will tend to act as 'attractors' in respect of less stable events, while formal approaches, such as those of Balzano (1982) and of Browne (1981), propose that certain formal properties of the pitch collections used in western common practice period music are more, or less, likely to motivate the use of particular configurations of pitches.

The work presented in the present study is most closely related to that of Balzano in that its starting point is the consideration of the implications for pitch cognition of group-theoretic properties manifested in canonical forms of western musical pitch organisation since the common-practice period. The canonical form in question is the cyclic group of order 12 which results from conceiving of pitches that are octave related as being equivalent, yielding the chroma circle and the circle of fifths (as isomorphic forms of the cyclic group). In a previous paper (Woolhouse & Cross, 2004) we outlined a theory of 'interval periodicity' (here referred to as ‘interval cycles’) within which relationships between pitches, and between pitches and sets of pitches, could be conceived of in terms of a simple additive model which took as its basic metric the different periods (cycles) manifested by different intervals between pitches within the isomorphs of the cyclic group. We presented empirical evidence demonstrating that the simple additive model was able to predict accurately the extent to which listeners would experience 'attraction' between dyads and subsequent single pitches. In the current paper the model previously described is extended and further empirical evidence for its predictive powers in the cognition of pitch attraction is presented.

MODEL

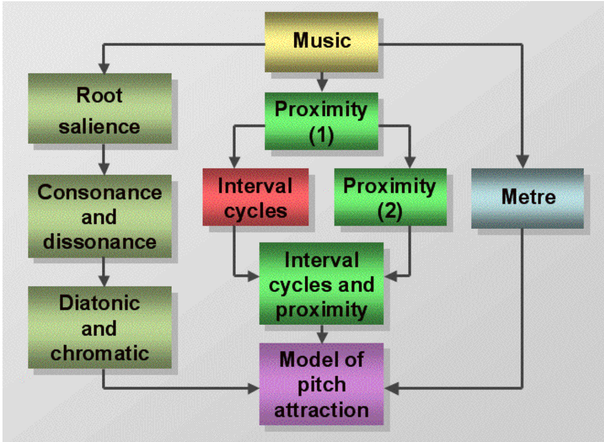

Figure 1. Interval cycle-based model of pitch attraction

The model has a number of components that may be switched ‘on’ or ‘off’ depending on the style of music being modeled. For example, if the attraction between functionally unambiguous diatonic chords is being modeled, *root salience* (a method giving the chords’ root relationships greater weight) is included. However, if the attraction between functionally ambiguous chromatic chords is being modeled, *root salience* is excluded.

The following is a detailed formal description of the components of the model. After each component, a short commentary and an illustrative example are provided. Component *Metre* is not covered in this paper.

Music

Given two successive chords (a chord may contain one or more notes): past chord, A, and present chord, B.

$A = \{a_1, a_2, \ldots, a_{|A|}\}$ $\qquad a_1 < a_2 < \ldots < a_{|A|}$

$B = \{b_1, b_2, \ldots, b_{|B|}\}$ $\qquad b_1 < b_2 < \ldots < b_{|B|}$

Where |X| = size of set X, and where a_i and b_j are defined with reference to $C_4 = 60$.

Defines the number of pitches in both chords and assigns each an integer value relative to C_4. Example, Vb to Ia in C major:

Vb (past chord) = $\{B_3, D_4, G_4\}$ = A = {59, 62, 67}

Ia (present chord) = $\{C_4, E_4, G_4\}$ = A = {60, 64, 67}

Proximity 1

Form matrix P where

$P_{ij} = |b_j - a_i|$ $\qquad i = 1, 2, \ldots, |A|$

$\qquad j = 1, 2, \ldots, |B|$

ISBN 88-7395-155-4 © 2006 ICMPC

Pitch proximity matrix in which absolute pitch distance is measured in semitones. Example:

Matrix P		B (Ia) 60	64	67
A (Vb)	59	1	5	8
	62	2	2	5
	67	7	3	0

Proximity 2

Define a further matrix Q such that

$$Q_{ij} = \frac{f}{P_{ij} + f} \quad \text{for some } f > 0$$

Scales matrix P by some value f. This limits and controls the influence of proximity. Example, f = 50:

Matrix P		B (Ia) 60	64	67
A (Vb)	59	1	5	8
	62	2	2	5
	67	7	3	0

→

Matrix Q (f = 50)		B (Ia) 60	64	67
A (Vb)	59	.98	.91	.86
	62	.96	.96	.91
	67	.88	.94	1.0

Interval cycles

Form matrix R where

$$R_{ij} = \frac{12}{hcf(P_{ij},\ 12)} \quad \text{for } P_{ij} \neq 0$$

$R_{ij} = 1$ if $P_{ij} = 0$

Where *hcf* (a, b) is the highest common factor of a and b, and a, b are whole numbers.

Matrix P proximity entries are converted into matrix R interval cycle entries. Example:

Matrix P		B (Ia) 60	64	67
A (Vb)	59	1	5	8
	62	2	2	5
	67	7	3	0

→

Matrix R		B (Ia) 60	64	67
A (Vb)	59	12	12	3
	62	6	6	12
	67	12	4	1

Interval cycles and proximity

Using entry-wise multiplication, combine matrix Q with matrix R, to give a further matrix S defined by

$S_{ij} = Q_{ij}\, R_{ij}$

Matrix Q (*Proximity 2*) is combined with matrix R (*Interval cycles*). The result is matrix S. Example:

Matrix Q (f = 50)		B (Ia) 60	64	67
A (Vb)	59	.98	.91	.86
	62	.96	.96	.91
	67	.88	.94	1.0

Matrix R		B (Ia) 60	64	67
A (Vb)	59	12	12	3
	62	6	6	12
	67	12	4	1

Matrix S		B (Ia) 60	64	67
A (Vb)	59	11.76	10.91	2.59
	62	5.77	5.77	10.91
	67	10.52	3.77	0.00

Root salience

Form matrix T where

$T_{ij} = 1 \qquad i = 1, 2, \ldots, |A|$

$\qquad\qquad\; j = 1, 2, \ldots, |B|$

If either chord has a clearly identifiable root, let the row corresponding to the past root be the m^{th} row and the column corresponding to the present root be the n^{th} column.

Form matrix U where

$$U_{ij} = \begin{cases} T_{ij} \text{ if } i \neq m \\ T_{ij} \text{ if } j \neq n \\ f \times T_{ij} \text{ if } i = m \text{ for some } f > 1 \\ g \times T_{ij} \text{ if } j = n \text{ for some } g > f \end{cases}$$

such that root intersection entry $U_{mn} = f \times g \times T_{mn}$. If neither chord has a clearly identifiable root, form matrix U where

$U_{ij} = T_{ij}$

Form matrix V where

$$V_{ij} = \frac{U_{ij}}{\sum_{i=1}^{|A|} \sum_{j=1}^{|B|} U_{ij}}$$

ISBN 88-7395-155-4

The idea behind *Root salience* is that in diatonic music root relationships are of greater perceptual importance. First, a new matrix, T, of identical dimensions to matrix P, is formed in which all entries = 1. Example:

Matrix T		B (Ia) (n) 60	64	67
A (Vb)	59	1	1	1
	62	1	1	1
	(m) 67	1	1	1

Second, the entries in row m (past chord's root) are multiplied by some number greater than 1, f. Third, the entries in column n (present chord's root) are multiplied by some number greater than f, g. This ensures that the root to root relationship of the chords has a greater weight than the root to non-root relationships, and that the present root has a greater weight than the past root (i.e. reflecting the time-dependent percept of present over past). It also ensures that non-root to non-root relationships have the least weighting, and by implication are of least perceptual importance. Example, f = 4, g = 8:

Matrix U		B (Ia) (*8) 60	64	67
A (Vb)	59	8	1	1
	62	8	1	1
	(*4) 67	32	4	4

→

Matrix V		B (Ia) 60	64	67
A (Vb)	59	.13	.02	.02
	62	.13	.02	.02
	67	.54	.07	.07

Matrix V is a normalization of matrix U, such that the sum of the entries in matrix V is equal to 1. Chromatic music frequently has no clearly identifiable root, and as a result does not produce different weightings in matrix U. Matrix V therefore equalizes diatonic music and chromatic music within the parameters of the model.

Consonant and Dissonant

(i) If the past and present chords are both sensory consonances or both are sensory dissonances, form matrix W where

$W_{ij} = V_{ij}$

Sensory consonance and dissonance are treated as time-dependent percepts. However, if both chords are consonant or both are dissonant matrix V remains unchanged.

(ii) If the past chord is dissonant and the present chord is consonant, form matrix W where

$W_{ij} = f \times V_{ij}$ for some $f > 1$

If the past chord is dissonant and the present chord consonant, i.e., there is temporal movement from dissonance to consonance, then the overall attraction between chords is increased.

(iii) If the past chord is consonant and present chord is dissonant, form matrix W where

$W_{ij} = V_{ij}/f$ for some $f > 1$

If the past chord is consonant and the present chord dissonant, i.e., there is temporal movement from consonance to dissonance, then the overall attraction between chords is reduced.

Diatonic and Chromatic

(i) If the past and present chords are both diatonic or both are chromatic form matrix X where

$X_{ij} = W_{ij}$

Diatonicism and chromaticism are treated as time-dependent percepts and are akin to the effect of consonance and dissonance. However, if both chords are diatonic or both are chromatic matrix V remains unchanged

(ii) If the past chord is chromatic and the present chord is diatonic, form matrix X where

$X_{ij} = f \times W_{ij}$ for some $f > 1$

If the past chord is chromatic and the present chord is diatonic, i.e., there is temporal movement from chromaticism to diatonicism, then the overall attraction between chords is increased.

(iii) If the past chord is diatonic and the present chord is chromatic, form matrix X where

$X_{ij} = W_{ij}/f$ for some $f > 1$

If the past chord is diatonic and the present chord is chromatic, i.e., there is temporal movement from diatonicism to chromaticism, then the overall attraction between chords is reduced.

Model of pitch attraction

Using entry-wise multiplication, combine matrix S with matrix X to give a further matrix Y defined by $Y_{ij} = S_{ij} X_{ij}$. Sum all entries in Y to produce global value G, where

ISBN 88-7395-155-4

$$G = \sum_{i=1}^{|A|}\sum_{j=1}^{|B|} Y_{ij}$$

Matrix Y combines interval cycle and proximity information from matrix S with root salience, consonance and dissonance, and diatonic and chromatic information from matrix X. Finally, the entries of matrix Y are summed to produce a single global total, G. Example:

Matrix S		B (Ia) (*8) 60	B (Ia) 64	B (Ia) 67
A (Vb)	59	11.76	10.91	2.59
	62	5.77	5.77	10.91
	(*4) 67	10.52	3.77	0.00

Matrix X		B (Ia) 60	B (Ia) 64	B (Ia) 67
A (Vb)	59	.13	.02	.02
	62	.13	.02	.02
	67	.54	.07	.07

Matrix Y		B (Ia) 60	B (Ia) 64	B (Ia) 67
A (Vb)	59	1.53	.22	.05
	62	.75	.12	.22
	67	5.68	.26	.00

G = 8.83 (the summed entries of Y)

MODEL PREDICTIONS 1

The model makes a number of readily testable predictions. For example, due to its importance in tonal music, of particular interest are the model's predictions for the attraction level of the C major triad in relation to the other diatonic triads d, e, F, G, a, b^o. From a music theoretic perspective, the C major triad – the *tonic* in the key of C – is strongly associated with G, F and b^o (Piston, 1941, p.21), due to their high transitional probability.

The first experiment aimed to test the hypothesis that the music-theoretic associations outlined above in relation to the C major triad are, in part, attributable to the attraction levels between chords as specified by the model.

EXPERIMENT 1: CHORD TO CHORD

Stimuli

The stimuli design was based on chord-pairs. For the major key, the triads C, d, e, F, G, a and b^o were paired with each other in both temporal orders. For example, in addition to C being paired with d in temporal order C (antecedent) followed by d (consequent), C was also paired with d in the reverse order, d (antecedent) followed by C (consequent). For the minor key, the triads c, d^o, Eb, f, G, Ab and b^o were paired.

Following the elimination of chord-pairs that are intervallic duplications of each another, for example G to C is transpositionally identical to C to F, 57 unique chord-pairs were identified. These chord-pairs formed the stimuli of the experiment.

Each chord was two seconds long, the consequent chord following the antecedent without a break. Each chord-pair was therefore four seconds long.

Method

Subjects

The subjects were 11 adults and included 5 females and 6 males with a mean age of 29 (SD = 8). All the subjects were members of the Faculty of Music, Cambridge University, UK. The subjects in the study were largely involved with western classical music, had received on average over 10 years formal musical training, and professed to practising/playing/performing on average 6.6 hours per week. The subjects' practical involvement with classical music was reflected in their listening habits which showed a 55% - 45% split in favour of classical music over popular/jazz/rock/folk music, and which was on average 7.4 hours per week. The subjects received no financial remuneration for their participation in the study.

Apparatus

Stimuli were produced using an Apple Macintosh G5 computer running *SuperCollider3* software. 'Shepard' tones (Shepard, 1964), consisting of superposed octave-related sinusoids with overall amplitude controlled by a Gaussian function, were used for the presentation of the stimuli. The technique produces tones with approximately equal overall pitch height and with no clear pitch maxima or minima. Subjects listened via headphones and adjusted the volume to a comfortable level before testing commenced.

Approximately the first 10 chord-pairs presented to each subject were practice stimuli and discounted from the analysis. To counter order effects, each subject was presented with the stimuli in a different random order and transposition.

Task

Subjects were required to rate on a seven-point scale (via a computer interface) the degree of attraction and/or resolution they felt between the chord-pairs: seven for a very strong sense of attraction, one for a very weak sense of attraction. In order not to confuse this task with the 'goodness of fit' probe tone paradigm used by, for example, Krumhansl & Kessler (1982), it was suggested to subjects prior to testing that a chord in relation to itself, i.e., chord repetition, should be considered as possessing a relatively low level of attraction/resolution. Following each chord-pair, subjects were given an unrestricted length of time in which to record their response. The subjects repeated the experiment and were therefore presented with each chord-pair twice during the session, from which a mean rating was taken.

Analysis

Prior to analysing the data, the stimuli presentation order was unscrambled and re-transposed in terms of C major. In order to equalize subjects' use of the seven-point scale, the results are presented as Z-scores.

The model's efficacy as a predictor of chord-pair attraction was assessed by calculating the Pearson correlation coefficient between the model and data.

Results

In order to test the central tenet of the theory – interval cycles – the data are modeled using as few components of the model as possible. The following correlations with the data are achieved using only two components: *Interval cycles* and *Root salience*. Components *Proximity 2*, *Consonance and dissonance*, and *Diatonic and chromatic* are excluded from the model.

Due to limited space this report will focus on the results obtained for the subset of chord-pairs pertaining to the C major triad.

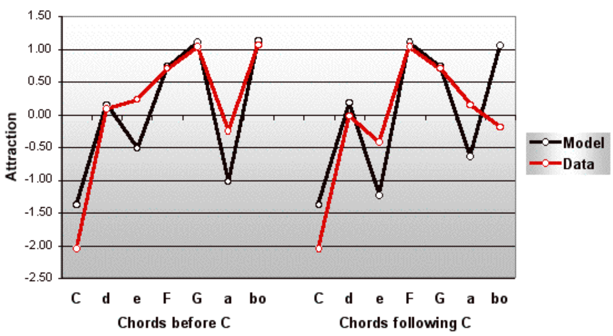

Figure 2. Graph showing the model's prediction of chord attraction (black) and subject's standardized mean responses (red) for the C major triad.

The left of Figure 2 shows the model and data attraction profiles for chords before C, i.e. d to C, e to C, F to C, and so on. The right of Figure 2 shows the model and data attraction profiles for chords following C, i.e. C to d, C to e, C to F, and so on. The correlation between the model and data is: r = .81; n = 14; $p < 0.005$. Excluding the C to C progressions on the grounds that for this progression the subjects were, to some degree, coached (see section *Task* above), the correlation falls to .75, which for n = 12 is nonetheless still highly significant ($p < 0.005$).

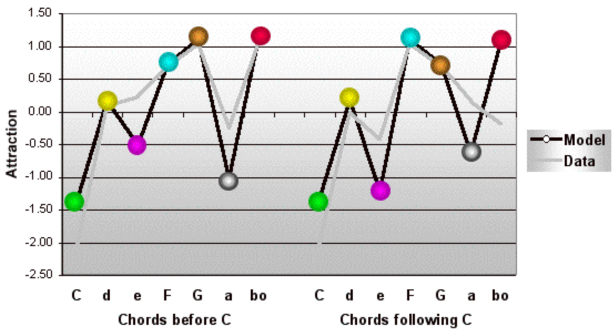

Figure 3. Graph showing the model's prediction of chord-pair perceptual asymmetries.

By colour-matching chord-pairs, i.e., d to C and C to d yellow, e to C and C to e purple, F to C and C to F blue, Figure 3 shows the model to predict a number of chord-pair perceptual asymmetries. That is, the attraction from, for example, e to C to be higher than from C to e, from F to C to be lower than C to F, and so on.

To show this more clearly, and to observe whether the subjects' data corresponded to these predictions, the values for the chord-pairs on the right of Figure 3 were subtracted from the corresponding chord-pair values on the left of Figure 3 for the model and subject data respectively. For example, the values for C to d were subtracted from d to C; the values for C to e were subtracted from e to C, and so on. The result is displayed in Figure 4.

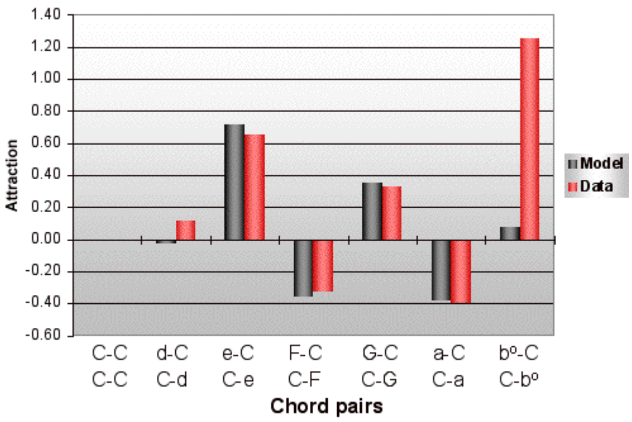

Figure 4. Model and data chord-pair subtraction. On the x-axis the bottom line of chord-pairs were subtracted from the top line of chord pairs.

Discussion

Figure 2 clearly shows the effectiveness of the model in predicting the attraction of the C major triad in relation to the other diatonic triads. Most noticeable on the left of Figure 2 are the high values in the model and data for the chord-pairs F to C, G to C and b^{o} to C. This result is in line with music-theoretic statements relating to the transitional probabilities of these chords, and offers a novel explication for their temporal association (Piston, 1941).

Two aspects of the profiles in Figure 2 are less successful, however: the mismatch between model and data for chord-pairs e to C and C to b°.

The higher subject rating for chord-pair e to C compared to the model can perhaps be accounted for by the leading tone progression B to C that the chord-pair e to C outlines. It would appear that in this instance, subjects focused their attentions on a relatively narrow aspect of the chord-pair transition – one traditionally associated with resolution – rather than on a global percept incorporating all aspects of the stimuli.

The lower subject rating for chord-pair C to b° compared to the model is indicative of the difficulty subjects had in experiencing a sense of attraction when the temporal order was from a sensory consonance to a sensory dissonance. This may, in part, be due to the terms in which the experiment task was framed, in that the concepts of 'attraction' and 'resolution' were conflated (see section *Task* above). Clearly the two expressions are related and not easily distinguishable. However, as a consequence, it is likely that a future experiment may take account of this and attempt, as far as it is possible, to define 'attraction' and 'resolution' unequivocally.

Despite the above problem, the component *Consonance and dissonance* (not included in the data analysis) can be used to model the sense of attraction when the temporal order is from a sensory consonance to a sensory dissonance. If the attraction from C to b° is recalculated using the model's *Consonance and dissonance* component, the Pearson correlation coefficient, excluding the chord-pair C to C, rises from .75 to .94 (n = 12; $p < 0.005$).

Finally, to a large extent the model correctly predicted the direction of the perceptual asymmetries (see Figure 4). For example, the model and data agree that the attraction of chord-pair e to C is greater than C to e, and that the attraction of chord-pair F to C (*plagal cadence)* is weaker than F to C (*perfect cadence*). Note too the relatively similar values given by the model and subjects to chord-pairs C to d and d to C. That is, the level of attraction from C to d and d to C was largely equal. For reasons relating to sensory consonance dissonance already referred to, the largest discrepancy between the model and data occurred for chord- pairs C to b° and b° to C.

MODEL PREDICTIONS 2

In order to model the attraction of functionally unambiguous diatonic triads that have, with the exception of b°, clearly defined roots, Experiment 1 used two components of the model, *Interval cycles* and *Root salience*. However, many chromatic chords – and to some extent, diatonic chords within chromatic contexts – are functionally ambiguous, with ill-defined roots. In terms of the model, this means that functionally ambiguous chromatic chords, in which no pitch necessarily carries greater perceptual weight, can in theory be modeled using only the central component of the model, *Interval cycles.*

Despite its often problematic functional description, chromatic music can nevertheless induce strong feelings of transitional pull and attraction, and is therefore – experientially at least – not dissimilar to goal-directed diatonic functional harmony. However, due to the difficulties of ascribing a definitive function to many chromatic chords, the directional pull – what might be referred to as a chord's 'attraction propensities' – are, from a music theoretic perspective, harder to predict and describe.

The second experiment aimed to test the attraction propensities of chromatic chords using a simple probe tone technique.

EXPERIMENT 2: CHROMATIC CHORD TO PROBE TONE

Stimuli

The stimuli design was based on a chord followed by a probe tone. The chords used in the study were the *dominant 7th* (or *German 6th*), the *half-diminished 7th* (or *Tristan chord*), and the *French 6th*.

The chord and probe tone were each two seconds long; the probe tone followed the chord without a break. Each chord plus probe tone was therefore four seconds long. Each chord was paired with probe tones transposed to all 12 possible pitch classes.

Method

Subjects and Apparatus

The subject group and stimuli production were as per Experiment 1.

Prior to testing each subject responded to approximately 10 practice stimuli. To counter order effects, each subject was presented with the test stimuli in a different random order and transposition.

Task

Subjects were required to rate on a seven-point scale (via a computer interface) the degree of attraction and/or resolution they felt from the chord to the probe tone: seven for a very strong sense of attraction, one for a very weak sense of attraction. Following each chord to probe tone pair, subjects were given an unrestricted length of time in which to record their response. The subjects repeated the experiment and were therefore presented with each chord/probe twice during the session, from which a mean rating was taken.

Analysis

Prior to analysing the data, the stimuli presentation order was unscrambled and re-transposed in terms of a single pitch. In order to equalize subjects' use of the seven-point scale, the results are presented as Z-scores.

The model's efficacy as a predictor of pitch attraction was assessed by calculating the Pearson correlation coefficient between the model and data.

Results

The following correlations between model and data are achieved using only one component: *Interval cycles*. Components *Root* salience, *Proximity 2*, *Consonance and dissonance*, and *Diatonic and chromatic* are excluded from the model.

*Dominant 7^th^*: G, B, D, F (or *German 6^th^*: G, B, D, E#)

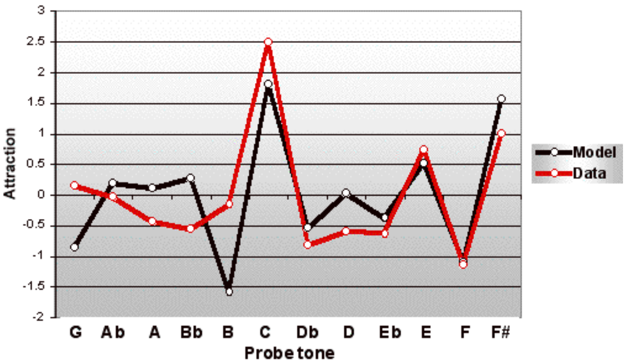

Figure 5. Graph showing the model's prediction (black) and subject's mean responses (red) for the attraction from the *dominant 7^th^* (G, B, D, F) to the probe tones.

The correlation between the model and data is: r = .76; n = 12; $p < 0.005$.

Half-diminished 7^th^ chord: D#, F, G#, B (or *Tristan chord*)

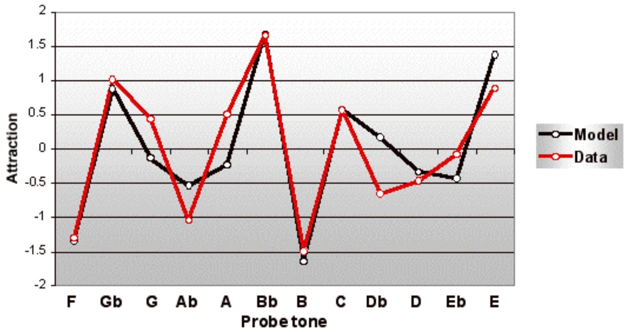

Figure 7. Graph showing the model's prediction and subject's mean responses for the attraction from the *half-diminished 7^th^ chord* (D#, F, G#, B) to the probe tones.

The correlation between the model and data is: r = .89; n = 12; $p < 0.005$.

*French 6^th^*: D#, F, A, B

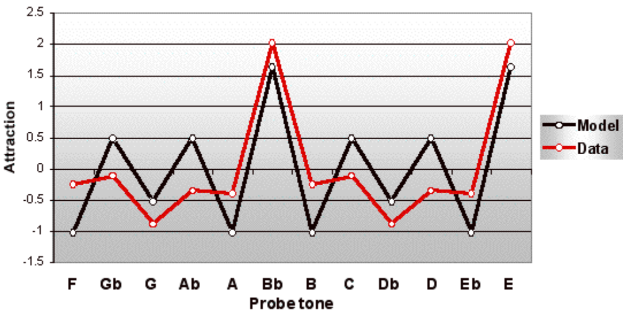

Figure 6. Graph showing the model's prediction and subject's mean responses for the attraction from the *French 6^th^* (D#, F, A, B) to the probe tones.

The correlation between the model and data is: r = .79; n = 12; $p < 0.005$.

Discussion

Figures 5, 6 and 7 clearly show the effectiveness of the model in predicting the attraction the chords to the probe tones.

The subjects' and model's attraction profiles for the *dominant 7^th^* shows two peaks: C and F#. While the subjects' rating peak on C following a G^7 chord was not unexpected, the rating peak on F# would seem to indicate that the *dominant 7th* was also interpreted chromatically, i.e., as a *German 6^th^* resolving downwards via a semitone step.

The correlation coefficients of the subjects' and model's attraction profiles for the *half-diminished 7^th^ chord* and *French 6^th^* were .89 and .79 respectively. Given the musically unorthodox nature of the stimuli and experiment task, these highly statistically significant correlations would appear to be a clear demonstration of the robustness of the model.

ISBN 88-7395-155-4

Conclusion

The *half-diminished 7th chord* and *French 6th* together form the central part of perhaps the most enigmatic and theoretically controversial progressions in western tonal music (Erickson, 1975): the opening of Wagner's *Tristan und Isolde*.

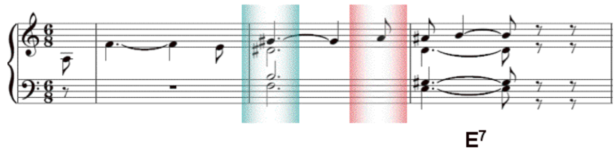

Figure 8. Richard Wagner, *Tristan und Isolde* (bars 1-3)

In Figure 8 the *Tristan chord* is highlighted in blue, the *French 6th* in red, the destination of the phrase is E^7. From the model and subject data the attraction towards the pitch forming the functional root of the chord at the end of the phrase, E, is strong (Figures 6 and 7) and therefore – despite the theoretical intractability of the harmony – it may be conjectured that the perceptual direction of the music is established within the *Tristan chord*, strengthened in the *French 6th*, and 'resolved' at the close of the phrase.

Clearly from the attraction profiles in Figures 6 and 7 a number of directional possibilities are present. That a model built solely on interval cycles can achieve significant correlations with the subjects' ratings given these complex stimuli is testimony to the strength of interval cycles in musical pitch cognition.

ACKNOWLEDGMENTS

The authors wish to thank Nick Collins for programming the experiments in *Supercollider3*. Matthew Woolhouse is generously supported by Arts & Humanities Research Council, UK.

REFERENCES

Balzano, G. J. (1982). The Pitch Set as a Level of Description for Studying Musical Pitch Perception. In M. Clynes (Ed.), *Music, Mind, and Brain*. New York: Plenum.

Bregman, A. S. (1991). *Auditory Scene Analysis*. Cambridge, Mass: MIT Press.

Browne, R. (1981). Tonal Implications of the Diatonic Set. *In Theory Only*, 5, 3-21.

Erickson, Robert (1975). *Sound Structure in Music*. Oakland, California: University of California Press.

Huron, D. (2001). Is music an evolutionary adaptation? *Biological Foundations of Music* (Vol. 930, pp. 43-61).

Krumhansl, C. L., & Kessler, E. J. (1982). Tracing the Dynamic Changes in Perceived Tonal Organization in a Spacial Representation of Musical Keys. *Psychological Review*, 89.

Krumhansl, C. L. (1996). A perceptual analysis of Mozart's Piano Sonata K.282: Segmentation, tension, and musical ideas. *Music Perception*, 13(3), 401-432.

Krumhansl, C. L. (2000). Tonality induction: A statistical approach applied cross- culturally. *Music Perception*, 17(4), 461-479.

Kurth, E. (1931). *Musikpsychologie.* Berlin.

Larson, S. (1998). Musical forces and melodic patterns. *Theory and Practice*, 22/23, 55-71.

Lerdahl, F., & Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press.

Lerdahl, F. (2001). *Tonal Pitch Space*. Oxford: Oxford University Press.

Oram, N., & Cuddy, L. L. (1995). Responsiveness of Western Adults to Pitch-Distributional Information in Melodic Sequences. *Psychological Research-Psychologische Forschung*, 57(2), 103-118.

Parncutt, R. (1989). *Harmony: A Psychoacoustical Approach*. Berlin: Springer-Verlag.

Piston, W. (1978). *Harmony* (4th ed.). New York: Norton. (Originally published, 1941.)

Shepard, R.N. (1964). Circularity in judgments of relative pitch. *Journal of the Acoustical Society of America*, 36, 2346-53.

Terhardt, E. (1974). Pitch, Consonance, and Harmony. *Journal of the Acoustical Society of America*, 55, 1061-1069.

Tillmann, B., Bharucha, J. J., & Bigand, E. (2000). Implicit learning of tonality: A self-organizing approach. *Psychological Review*, 107(4), 885-913.

Woolhouse, M.H., & Cross, I. (2004). The role of interval periodicity in hierarchical pitch perception. *Eighth International Conference on Music Perception and Cognition,* 486-492.