Computer-assisted composition with algorithmic implementations of pitch-perceptual theory

Richard Parncutt (University of Graz) and Sean Ferguson (McGill University)

This text was translated into Italian for publication as:

Parncutt, R., & Ferguson, S. (2005). Comporre con l'apporto di un algoritmo di teoria della percezione. Rivista di Analisi e Teoria Musicale, 10(2) (special issue "Musica e ricerca interdisciplinare" edited by Mario Baroni), 103-118.

The special issue of RATM was devoted to papers presented at the Conference on Interdisciplinary Musicology (CIM04), Graz/Austria, 15-18 April 2004

Date of this draft: 17.1.05

Abstract

We describe a computer-assisted procedure for composing harmonic progressions that takes into account the perceptual processes underlying roughness (sensory dissonance) and the pitch of complex tones. First, the user specifies the outer voices and specific model parameters, including preferred values of the perceived roughness of each sonority, the pitch commonality of successive sonorities, and the pitch commonality of each sonority with a reference sonority. The model then generates different possible chords and progressions within these constraints. The user is free to change the constraints to limit the number of possible chords and to choose from the generated chords according to personal criteria. The compositional process thus involves repeated interaction between composer and computer. We offer examples of progressions composed in this way from *Dans la chair/In the Flesh*, a recent piece for orchestra by the second author.

Introduction

High-speed computers have significantly affected most aspects of human culture, and musical composition is no exception. Computers enable the realisation and testing of complex mathematical compositional theories. Promising recent examples include the evolutionary approach of Manzolli, Moroni, Zubens, & Gudwin (1999) and the set-class approach of Castrén (2000). Both these approaches focus on the musical surface: How are individual sonorities perceived? How are harmonic relationships between successive sonorities perceived?

We ask much the same questions, but offer a different approach. Our theory is more perceptual or psychological than the theories of Zuben and of Castrén. We begin with a theory of pitch perception that was primarily intended to explain the perception of pitch in non-musical sounds, but also explains some interesting things about pitch in Western music. The only musical construct assumed in our computational approach is the chromatic scale, and even that is not central, since the program could easily be rewritten to generate music in other scales. The theory makes no prior assumptions about the consonance of musical intervals, chords, or chord progressions; instead, it makes assumptions about the way individual harmonic complex tones and sequences of sonorities are perceived, regardless of whether they appear in a tonal context. We will describe these assumptions in more detail below.

A further feature of our approach is that tones in different octaves are not treated as equivalent. A common objection to pitch-class theory is that a single pitch-class set can sound very different when realised in different voicings (inversions, spacings, doublings, registers). To avoid this problem, we retain octave register information throughout. We also retain information on the frequencies and the amplitudes of all audible spectral components.

We begin by exploring relevant scientific knowledge from the fields of auditory psychophysics and psychology. Researchers in these fields explore and measure *auditory capabilities* such as perceptual thresholds. They also investigate those *habits of perception* that are acquired as humans interact repeatedly with patterns of sound. In this paper we will be concerned with both. For example, not all of the frequencies that are physically present in a sound can be heard, because the ear's ability to separate nearby frequencies is limited.

Sound patterns in the human environment may be divided culture-specific and universal patterns. Culture-specific patterns may further be divided into patterns typical of speech and music. The habits of perception that develop when sound patterns are repeatedly heard can be divided into similar categories. For example, the well-known Gestalt principles of grouping by similarity, proximity, good continuation and so on are examples of apparently universal habits of perception, acquired in interaction with everyday visual and auditory stimuli including speech.

The formation of perceptual habits is an example of *auditory learning*. Learning may occur phylogenetically, that is, over several generations according to principle of survival of the fittest, with genetic/physiological transmission from one generation to the next. But the word "learning" is more properly applied to ontogenetic acquisition of information, that is, learning within a single lifetime. Universal habits of perception are not necessarily determined by the physiology of audition and cognition; they may also be learned from exposure to speech, since the physiology and acoustics of speech production are universal aspects. An example of a universal habit of auditory perception is the recognition of harmonic pitch patterns within complex tones, which is either phylogenetically or ontogenetically related to the periodicity, and hence harmonicity, of individual voiced speech sounds. This habit of perception plays an important role in the theory that we will develop below.

In the past century or so, scientists have investigated a wide variety of auditory capabilities and habits of perception. This body of knowledge would appear to have important implications for the much older art of musical composition, and raises the following question: To what extent do, or should, *compositional structures* reflect the auditory capabilities and habits of *listeners*?[1]

For nineteenth-century music theorists and psychologists such as Helmholtz, Kurth, Riemann and Stumpf, the answer to the above question may have been so clear that there was no need to discuss it. Compositional structures obviously reflect the auditory capabilities and habits of most people, otherwise the music would be incomprehensible. Their question was rather: What auditory capabilities are relevant to music perception, and what is the nature of the correspondence between these capabilities are musical structures? Questions of this kind were pursued in the early 20th century by the Gestalt psychologists.

Musically and artistically, the late 19th century was characterized by an accelerating drive toward originality. The romantic ideal of unfettered expression of individual creative imagination and the resultant breaking down of classical aesthetic traditions of content, convention and genre was accompanied by an expansion of the tonal system towards its limits and, eventually, beyond them. By the time academic musicologists finally accepted the expanded harmonic resources of the early decades of twentieth century into the musical canon, modernist twentieth century composers and musicologists were freely and openly rejecting the notion that composition should be limited by perceptual constraints of any kind. For example, Babbit claimed that because composers are experts, listening to their music requires expert listening (Peles, 2003). On this basis, and driven by the overriding goal of originality, modernist composers believe(d) that new music could overcome any perceptual constraint: no matter how advanced the musical hearing of a given listener or musical subculture, new music could help them to open their ears to an even broader range of musical structures and to even less "natural" soundscapes. This attitude contributed to the 20th century's unprecedented stylistic richness and diversity.

---

[1] By “compositional structures”, we primarily mean patterns of pitch against time, such as the specific voice-leading generated by a cadence. More specifically, we are referring to patterns involving *notatable perceptual categories* such as scale steps (pitch categories) and note values (rhythmic categories). The definition could be extended to include patterns of timbres in the sense of physical sources of sound, which may also be regarded as notatable perceptual categories - but not to spectromorphology. By “listeners”, we mean Western listeners, since the theories we invoke may only apply within Western culture. Although perceptual capabilities and habits clearly depend on musical tastes and expertise, the theory that we will apply is independent of musical training, so our listeners could be either musicians or non-musicians.

The idea that new music can overcome any perceptual constraint was called into question by the international revival of music psychology toward the end of the twentieth century. For example, it has become clear that the human auditory system has clear preferences for tones in the central musical pitch range, which corresponds to the normal pitch range of human speaking and singing voices (Huron, 2001). Human listeners prefer melodies that move mainly by step, and identify them mainly according to their pitch contours (Dowling, 1978). Listeners tend to prefer structures that are tonal in the context of their own culture (in the west: major-minor), and preferences for tonality are associated with deep, complex, cognitive-structural constraints (Krumhansl, 1990). These findings do not mean that composers should be *restricted* by perceptual constraints; but they do imply that the perception of a composer's music will inevitably be *influenced* by perceptual constraints, and consequently that awareness of perceptual constraints on the part of the composer could enhance the compositional process and the chances that the resultant music will be accepted and valued by performers and audiences.

In incorporating perceptual theories into the compositional process, we are not attempting to eliminate aurally guided trial and error from the compositional process. Instead, we aim to complement and support traditional procedures of aural trial and error. We do this in two ways: by limiting the number of sounds from which a composer chooses by ear, and by pointing up new possibilities that the composer may not have thought of. For a simple, analogous example, consider the chords that typically follow a dominant seventh chord in tonal music. Tonal composers typically choose from among the limited range of possibilities that exist in previous music or in music theory, rather than from among all possible chords within the tonal system.

We take as our starting point Terhardt's (1974) and Parncutt's (1989) theories of the perception of tonal sonorities and of relationships between successive sonorities. In a nutshell, these authors proposed the following. A tonal sonority is made up of various pure-tone components (sine waves) with different frequencies and sound pressure levels. The perception of a tonal sonority involves two separate stages. First, mutual masking among the components reduces their individual audibility, and some may become completely inaudible. Second, harmonic patterns of audible components tend to perceptually fuse, that is, to be perceived as single complex tones; Terhardt called their pitches "virtual" and developed a model to predict their perceptual importance or *salience*. Parncutt evaluated the perceived relationship between successive sonorities by modeling the extent to which they had perceived pitches in common (pitch commonality). Here, the word *pitch* refers neither to the musical name of a note (e.g. C#) nor to the frequency of a pure or complex tone, but exclusively to the pitch of a tone as it is *experienced* – how high it sounds, both absolutely and relative to other tones. These distinctions play ~~in~~ an important role in this paper.

Hindemith (1937) had already suggested how such models might be applied in composition. He proposed controlling the dissonance of the chords in a progression, so that dissonance gradually increases towards the middle of a phrase and decreases towards the end. With the advent of personal computers, composers such as Barlow (1987) attempted to realise Hindemith's idea with computer-implemented compositional algorithms that evaluate musical or perceptual parameters such as dissonance or complexity.

Parncutt and Strasburger (1994) explained how modern pitch-perceptual theory could be incorporated into computer-based composition. They suggested first enumerating all possible sonorities within a given set of constraints, some of which relate to pitch perception (e.g., restricting the calculated salience of the most salient pitch in the sonority to a given range), and then enumerating all possible progressions of those sonorities within another set of constraints (e.g., restricting the pitch commonality of, and pitch distance between, successive pairs of sonorities, or of each sonority in relationship to a reference sonority, to given ranges).

Ferguson (2000) was the first to extensively and successfully realise this concept. First, he composed the outer voices of a passage. Then his computer model generated the inner voices of each chord in the progression, corresponding to predetermined, perceptually based parameters. As he selected the first chord from the possibilities offered by the computer, the computer offered possibilities for the second chord, and so on. In the process, he also adjusted the model parameters, or set them to change in a particular way during the course of a progression, until the computer output corresponded to his expectations or produced results that he considered interesting. Finally, he incorporated selected progressions into his composition. The software was written in the Common Lisp programming language. Note that this process of composition involves repeated *interaction* between computer and composer: the composer adjusts the parameters and outer voices, the computer offers possible musical materials, the composer selects from these by subjective and aesthetic evaluation, this selection limits the next round of computer predictions, and so on.

Before moving to the details of this compositional process, we should explain further what we mean by *perception of musical structure*. There is no clear answer to the question of whether a “structure” is “perceptible”: for example, one can “perceive” a melody as such, without having any idea of its structure. A cognitive psychologist might instead ask whether the cognitive representation of the melody (i.e. the mental structure that is assumed to exist in the mind of the listener when the melody is heard) corresponds to a given music-theoretical structure. We are not asking this question, either. Instead, we are assuming that, when music is composed according to empirically validated, general principles of pitch perception, the resultant musical structures will be more clearly perceptible and therefore make more sense - at least to those Western listeners for and with whom the theory was developed.

Underlying these various ideas is an aesthetic assumption. We are evidently assuming that music whose structures are “perceptible” is in some essential sense “better”, or more acceptable as “music”. When considering this claim it is important to remember that perceptibility depends on both structures and listeners. Thus, the same sounds can be regarded as non-music in one period and as music in another. Another point is that a composer who deliberately tries to avoid structure of any kind may, in so doing, create a new kind of structure that can be perceived on a different level.

Why do some musical styles survive long after they are composed, while others are forgotten? A possible reason is that the music’s sound structures continue to *sound interesting,* even as the musical and cultural context in which they are heard changes. This is only possible if the music’s sound structures are *perceptible*. If not, the music is reduced to an intellectual exercise.

Music examples in this paper are taken from *Dans la chair/In the Flesh*, a recent piece for orchestra by the second author that was premiered in Paris by the Orchestre Philharmonique de Radio-France during the Presences 2003 Festival. This music is quite complex and somewhat dissonant. The same perceptual principles and compositional procedures could be used to compose music that is less complex and more consonant, and appeals to a broader audience - but that would be the topic of a separate paper.

Harmonic vocabulary

The harmonic vocabulary of *Dans la chair/In the Flesh* was initially defined as a three-dimensional array. The dimensions were cardinality (number of pitch classes in each sonority), outer interval (interval between highest and lowest voice), and specific arrangement of pitches between those voices.

Chords were included in the array that correspond to a certain criterion, specified in terms of standard pitch-class set theory (Forte, 1973). That criterion was high similarity to a particular nexus set according to Castrén (1994). The chosen set was [0 1 4 6], one of the all-interval tetrachords. The final array contained a total of 47,603 chords with cardinalities from 3 to 10 and outer intervals up to 47 semitones.

The goal of restricting the harmonic vocabulary at the outset was twofold. First, the set of all possible chords with a given cardinality and outer interval can become extremely large. For example, while the number of all possible 6-note chords with an outer interval of 8 semitones is only 35, there are 2,380 possible 6-note chords with an outer interval of 18 semitones. Such large numbers are both musically unnecessary and computationally inefficient. Second, limiting the harmonic vocabulary to related chords may help to increase the coherence of a piece.

Psychoacoustic models

The compositional goal of the system was to create a harmonic syntax in which the following parameters are under the control of the composer: the roughness of chords, the relationship between successive chords in a progression, and the relationship of chords to a referential sonority (conceived of as a harmonic region). The two principal psychoacoustic models used were roughess and pitch commonality. Our working hypothesis was that, by basing syntax on perceptual models, a composer can control the perceived *harmonic stability* of a progression.

Roughness

The musical term *musical dissonance* has various meanings, of which we focus on two. The first, *stylistic dissonance*, is learned by exposure to music in a given style. It is connected to the cultural background and learning processes of the listener, and varies according to the historical style of the music (Terhardt, 1984; Tenney, 1988). Hutchinson and Knopoff (1979, p. 21) referred to "stylistic instability, tension, 'disagreement'". The modernist tradition of musical composition regards stylistic dissonance as the *only* form of musical dissonance; and since it is learned, it can also be unlearned. But research in music psychology has repeatedly demonstrated that musical dissonance also has sensory aspects. *Sensory dissonance* refers to any aspect of dissonance whose origin lies in universal properties of the auditory system (Plomp & Levelt, 1965) and may be considered independent of culture or musical style. The most robust such aspect is the sensation of rapid beating, or roughness, within a sonority. It is directly connected to the physiology of auditory frequency analysis on the basilar membrane in the cochlea. It is the main reason why a harmonic perfect fifth interval in the central musical pitch range (e.g. C4-G4) sounds relatively smooth, while a harmonic minor second (C4-Db4) sounds rough. Other factors that contribute to sensory dissonance are *noisiness* (as opposed to tonal clarity or pitch strength), *sharpness* (energy at high frequencies), and *loudness*.

Stylistic and sensory aspects of musical consonance and dissonance do not necessarily agree. The interval of a perfect 4th above a bass tone, for example, is considered to be dissonant or unstable in common-practice tonal syntax. This is an example of stylistic dissonance; the sensory dissonance of a harmonic perfect 4th interval is not much greater than that of a perfect fifth, when spectral envelopes and register are held constant. The minor 2nd, on the other hand, is dissonant according to both definitions. Stylistic factors also influence the way sensory factors are appraised: the harmonic major third interval was considered dissonant in early polyphony, but in the tonal music of the 17th-19th centuries it is a consonance.

It is important also to distinguish between *perceptual universals* and *musical universals*. We use the term "roughness" in the sense of a perceptual universal, because it arises primarily from the physiology of the ear. But the tendency in Western music to avoid roughness or to treat rough sonorities in particular ways is clearly not universal. Vocal duets from the Balkans sometimes deliberately emphasize harmonic second intervals (both major and minor) *because of* their roughness, and the aggressive sound of very loud or deliberately distorted rock music is in part due to roughness. The perceptual universal in this case is not the aesthetic appraisal of roughness, but the roughness itself. For purely physiological reasons, all people from all cultures who hear a harmonic second interval between harmonic complex tones of about the same intensity in the central musical pitch range perceive the dyad as rough – although they may use different words to describe the sensation, depending on their language and cognitive style. But different cultural traditions differ in how roughness is aesthetically appraised.

Helmholtz (1877/1954) investigated the relationship between beating and roughness experimentally. Two pure tones with almost the same frequency generate beats at a frequency equal to the difference between the two original frequencies. If this beat or difference frequency exceeds about 20 Hz, the individual beats become imperceptible, and the sound becomes rough. Plomp and Levelt (1965) proposed that roughness is perceived only when the beat or difference frequency is less than a *critical bandwidth*. Critical bandwidth is determined by the underlying physiology: each hair cell on the basilar membrane may be regarded as an acoustic filter that passes frequencies within a certain range (a *bandpass filter*). Critical bandwidth is the effective[2] width of that filter, measured in cycles per second (Hz) or semitones.

[2] The word "effective" is necessary because the filter is not rectangular. In other words, the boundaries of "critical band" are not exact or sudden. Instead, frequencies at greater distances from the centre are increasingly attenuated.

The roughness model included in our compositional algorithms is based on Hutchinson and Knopoff (1978), which was in turn based on Plomp and Levelt (1965).[3] Each pair of pure tones in the spectrum of a sonority is assumed to contribute to the sonority's roughness if the interval between them is smaller than a critical band. The overall roughness of the sonority is predicted by adding all such contributions.

Figure 1 lists all possible 4-note chords in the chromatic scale with an outer interval of a major 6th above middle C, and their roughness. The chords are arranged in order of increasing roughness according to the algorithm of Hutchinson and Knopoff. Each tone in each chord is assumed to contain 10 harmonic partials. The effectiveness of the model may be judged by playing the progression on the piano. Does the ranking correspond to your perception?

Pitch commonality

In common-practice tonal syntax, the number of common notes between successive chords is related to the strength of a progression: “weak” or “smooth” progressions have more common notes, and “strong" or “uneven” progressions have fewer (Aldwell & Schachter, 1989, p. 136). In pc-set theory, successive pc-sets sound smoother or more continuous when they have more pcs in common (termed “invariance”): "the concept of invariance is intimately bound up with the intuitive musical notions of development, change, continuity, and discontinuity" (Forte, 1973, p. 30).

Both these approaches to chord succession are restricted to the *notated pitches* of the chords. Another possibility is to consider common *frequencies*, as Helmholtz did in his theory of the harmonic affinity between successive sounds. Helmholtz's idea breaks down when a harmonic complex tone with missing fundamental is compared with a pure tone corresponding to that fundamental: the two are perceived as related although they have no frequencies in common. From this example, it is clear that a model of harmonic relationship should be based on *perceived* (or *experienced*) pitches, as pitch commonality is.

A mathematical model of pitch commonality was developed by the first author, based on Terhardt's (1974) revision of Helmholtz's concept of tonal affinity. The pitch commonality of two sonorities may be said to depend on their *perceptual spectra*. Whereas a regular physical (amplitude) spectrum is a graph of SPL against frequency for each partial, a perceptual spectrum is a graph of perceptual salience against the perceived pitch of each tone sensation. Pitch commonality is more complex to calculate than roughness; for a complete mathematical description of the model, see chapter 4 of Parncutt (1989), or Parncutt and Strasburger (1994).

In *Dans la chair/In the Flesh*, pitch commonality is applied in two ways: successively and referentially.

- Successive pitch commonality is calculated between successive chords in a harmonic progression. Figure 2 shows two different chord progressions with the same outer voices. In the first harmonization, successive chords have high pitch commonality, while the second version has low successive pitch commonality. The

[3] Our implementation of Hutchinson and Knopoff is published at www-gewi.uni-graz.at/muwi/parncutt.

pitch commonality between each pair of chords is given between the staves. These values are higher in the first progression, consistent with its smoother sound.

- Referential pitch commonality is calculated between the chords of a progression and a middleground referential sonority, which may either be present in the music, as a sustained chord for example, or merely suggested. Just as a major-minor tonal passage can imply the tonic before it is sounded, a progression of sonorities whose pitch commonality with a given sonority is consistently high may imply that sonority before if it is sounded. This effect may be stronger if the sonority has been previously emphasized in the piece.

The eight-note chord in Figure 3 is the referential sonority of *Dans la chair/In the Flesh*. The four-note chords to the right are sorted by decreasing pitch commonality with this sonority. Recall that pitch commonality involves *perceived* (rather than notated) pitches, and takes into account calculated variations in pitch *salience*. Recall also that these chords are defined as specific voicings – not of pitch-class sets. The maximum and minimum pitch commonality values are given below the staff. Chords at the beginning of the list sound more related to the referential sonority, since they share more pitches with it.

Composing progressions

We compose chord progressions by harmonizing an outer-voice framework. Inner voices are added to each pair of outer voices according to specific criteria. These include cardinality (number of notes), roughness, successive pitch commonality, and referential pitch commonality.

The first step is to find all the chords with the desired cardinality, outer interval, and transposition level. These are selected from the specific harmonic vocabulary of the piece as follows. The perceptual parameters take values between 0 and 1, where 0 is the smallest possible value of the parameter and 1 the largest, relative to currently available possibilities. For instance, a value of 0 for roughness tells the computer to find the least rough of all chords with the desired number of notes, outer interval and transposition level. A value of 0.5 means: "Of these chords, which ones have an average amount of roughness?" To answer this question, we sort the list by increasing roughness and retain only that segment of the list that corresponds to the desired amount. Very high and low values of the parameters tend to return fewer chords than intermediate values. The approximate portions of the list returned for several different values of roughness are given in Figure 4.

This procedure is repeated for each parameter in turn. The result of each successive "slice" is analyzed and sorted, and the corresponding portion is used for the next parameter. The order of analysis is: (1) roughness, (2) referential pitch commonality and (3) successive pitch commonality.

The following example clarifies the harmonization process. Imagine that we are harmonizing the fourth chord out of a progression of eighth chords, and that at this point the current values of the harmonic parameters are:

- referential sonority: (60 61 64 66) or C4, C#4, E4, F#4

- cardinality: 6
- transposition level: 61 (C#4)
- outer interval: 23 semitones
- roughness: 0.33
- referential pitch commonality: 0.67
- successive pitch commonality: 0.75

The program returns four progressively shorter lists of chords, notated as MIDI key numbers (middle C = 60). The first list contains all the chords in the harmonic vocabulary of *Dans la chair/In the Flesh* with the desired cardinality and outer interval, transposed to begin on MIDI note 61. The second list contains those chords from the first list whose relative roughness corresponds to the indicated amount of 0.33. The third list is the result of choosing chords from the previous list that have the desired degree of referential pitch commonality, and the fourth list is the result of analyzing these chords for successive pitch commonality. The composer makes a final choice from this list. Depending on the musical goals of the passage, one could choose the chord which is least rough, or more closely related to the referential sonority. Or, since the chords of the list are essentially equivalent in terms of the specified harmonic parameters, one can simply make a random choice. For the above progression, the four lists (in musical notation) are given in Figures 5a to 5d.

Example

Figure 6 shows a reduction of measures 45-49 of *Dans la chair/In the Flesh*. The bottom two staves contain the referential sonority of the piece, which is sustained in a complex texture which has been simplified for this example. The upper two staves contain harmonic progressions whose outer voices have been harmonized using the system described above.

Both progressions begin with high roughness, low successive pitch commonality and low referential pitch commonality, which in the course of the progression is gradually transformed to low roughness, high successive pitch commonality and high referential pitch commonality. The compositional goal is to create musical phrases with low harmonic stability at their beginnings, but which become gradually more harmonically stable as they progress.

Discussion

We have shown how perceptual theory can be used to create pieces and compositional styles that are complex and tonally original. We have presented a summary of the perceptual theory and computer programming, and discussed some examples.

Our approach is not intended to replace other similar compositional approaches, but to complement them. It expands the range of possibilities available to computationally agile composers. We hope with this article not only to have familiarized composers with our procedures, but also to have inspired the development of new hybrid procedures.

It is still unclear whether, or to what extent, the perceptual models incorporated into our compositional procedures might contribute to the music's success. It is possible that quite different and arguably less perceptibly logical pitch patterns might have achieved a similar effect. We are currently planning listening experiments to investigate this question. Sound materials for the experiment will be generated with the same software that supported the composition of *Dans la chair/In the Flesh*. Unlike this piece, which is entirely non-triadic and tonally so complex that it sounds atonal, the music examples in our experiments will cover a wide range from "very tonal" to "very atonal". By exploring quantitative relationships between aesthetic judgments and model parameter settings as well as qualitative analysis of listeners' verbal descriptions, we hope to shed more light on the general question of the relevance of perceptual psychology for composition.

References

Aldwell E., & Schachter, C. (1989). *Harmony and voice leading* (2nd. ed.). San Diego: Harcourt Brace Jovanovich.

Barlow, C. (1987). Two essays on theory. *Computer Music Journal, 11* (1), 44-60.

Castrén, M. (1994). *Recrel: A similarity measure for set-classes.* Ph.D. diss., Sibelius Academy, Helsinki.

Castrén, M. (2000). Pairs of chords as objects illuminating set-class similarity: Some viewpoints and a computer-assisted procedure to create test materials for listener experiments. *Electronic Journal of Music Theory and Analysis, 1* (1), **http://www.musictheoryresources.com/members/MTA_1_1a.htm.**

Dowling, W.J. (1978). Scale and contour: Two components of a theory of memory for melodies. *Psychological Review*, *85*, 341-354.

Ferguson, S. (2000). *Concerto for piano and orchestra*. Doctoral thesis, McGill University, Montreal, Canada.

Forte, A. (1973). *The structure of atonal music*. Yale University Press, New Haven.

Helmholtz, H.L.F. (1877/1954). *On the sensations of tone as a physiological basis for the theory of music* (transl. A. J. Ellis). Dover, New York.

Hindemith, P. (1937). *Unterweisung im Tonsatz*. Mainz: Schott.

Huron, D. (2001). Tone and voice: A derivation of the rules of voice-leading from perceptual principles. *Music Perception, 19*, 1-64.

Hutchinson, W., & Knopoff, L. (1978). The acoustic component of Western consonance. *Interface, 7,* 1-29.

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

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

Parncutt, R., & Strasburger, H. (1994). Applying psychoacoustics in composition: "Harmonic" progressions of "non-harmonic" sonorities. *Perspectives of New Music, 32* (2), 1-42.

Peles, S. (Ed.) (2003). *The collected essays of Milton Babbitt*. Princeton, New Jersey: Princeton University Press.

Plomp, R., & Levelt, W. J. M. (1965). Tonal consonance and critical bandwidth. *Journal of the Acoustical Society of America*, *38*, 548-560.

Tenney, J. (1988). *A history of 'consonance' and 'dissonance'.* New York: Excelsior.

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

Terhardt, E. (1984). The concept of musical consonance: A link between music and psychoacoustics. *Music Perception, 1*, 276-295.

Manzolli, J., Moroni, A., Zubens, F. von, & Gudwin, R. (1999). An evolutionary approach to algorithmic composition. *Organized Sound: An International Journal of Music Technology, 4* (2), 121-125.

Figure 1: Chords sorted by increasing sensory dissonance.

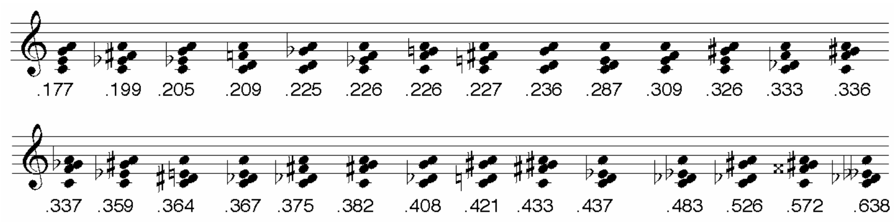

Figure 2: Two different harmonizations of the same outer voices.

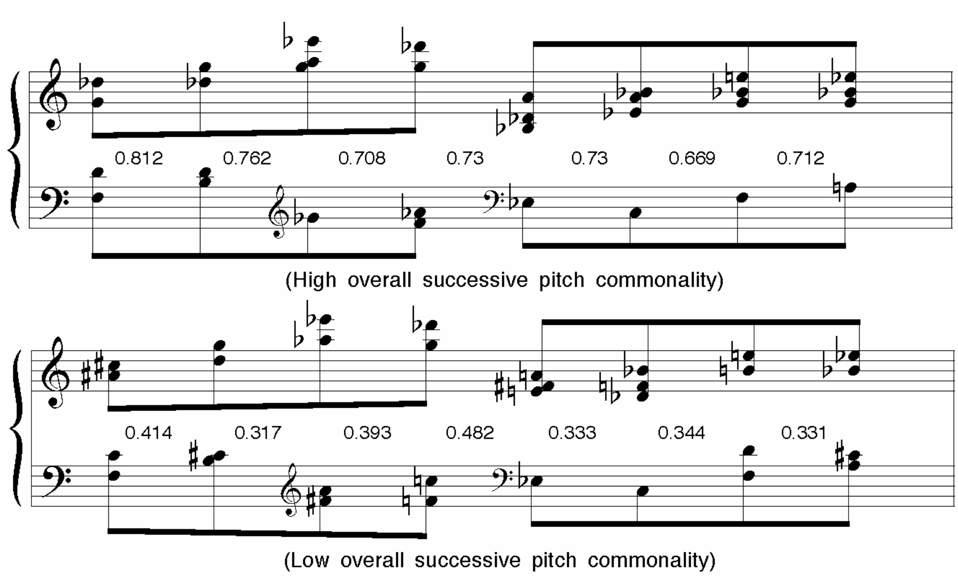

Figure 3: Chords sorted by referential pitch commonality.

Figure 4: Selecting chords by degree of sensory dissonance.

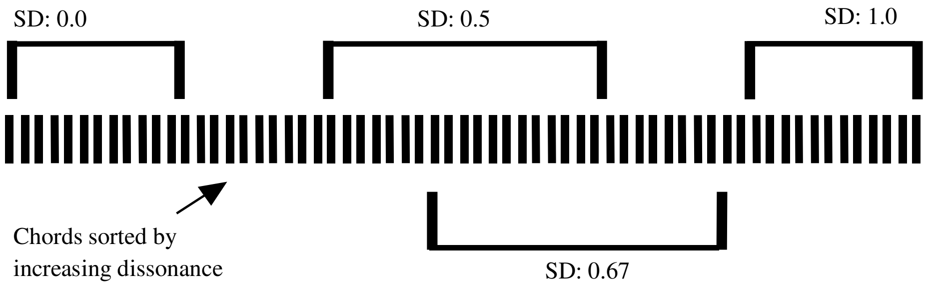

Figure 4: Selecting Chords by degree of sensory dissonance.

Figure 5: Examples of progressively shorter lists of chords output by the model.

Figure 5a: All Chords

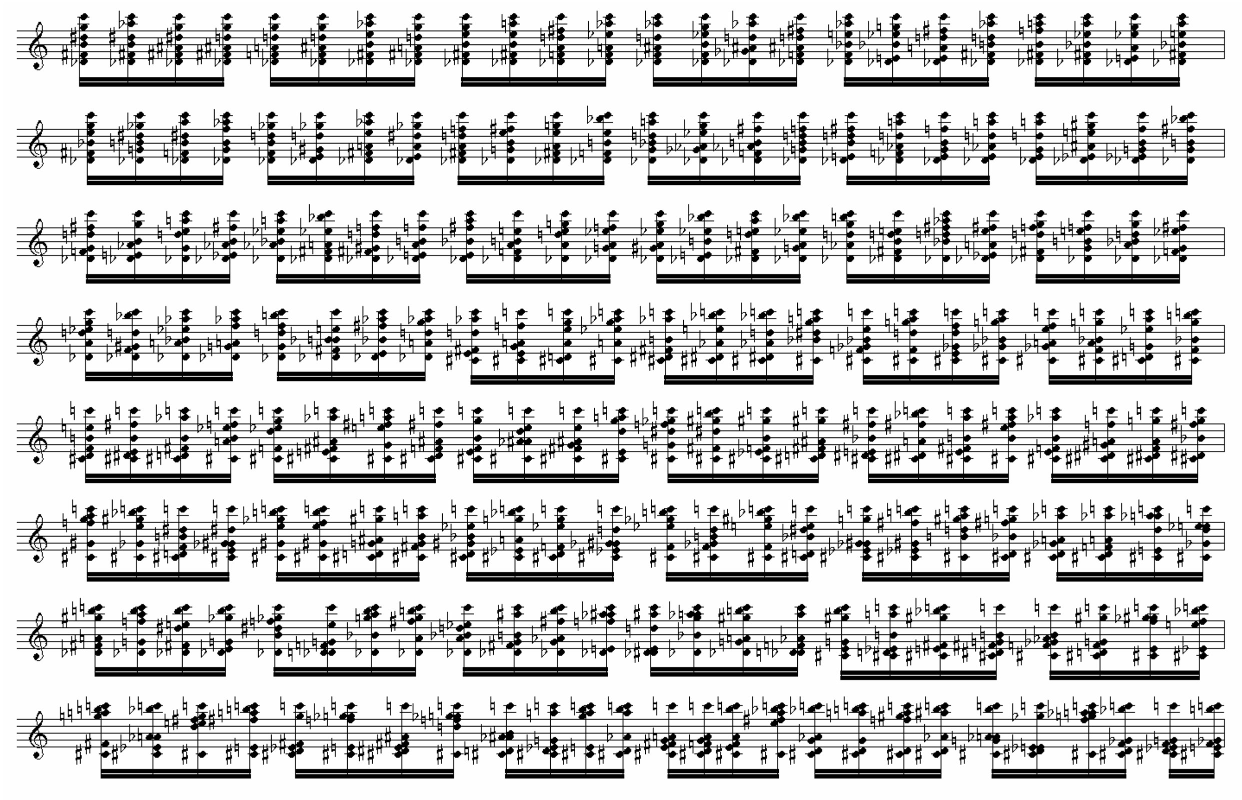

Figure 5b: Chords Selected By Sensory Dissonance

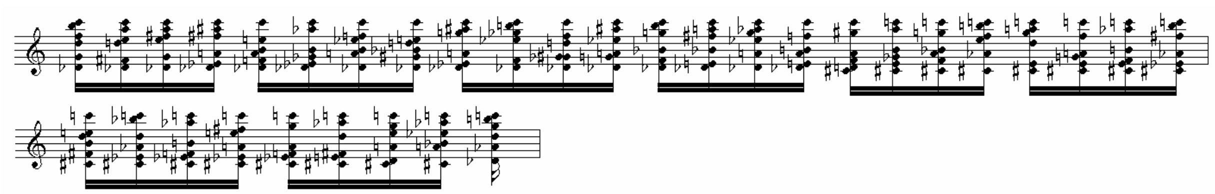

Figure 5c: Chords Selected By Referential Pich Commonality

Figure 5d: Chords Selected By Successive Pitch Commonality

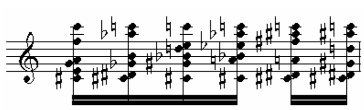

Figure 6: Measure 45-49 of *Dans la chair/In the Flesh.*