Conchord: An Application for Generating Musical Harmony by Navigating in the Tonal Interval Space

Gilberto Bernardes[1(✉)], Diogo Cocharro[1], Carlos Guedes[1,2], and Matthew E.P. Davies[1]

[1] Sound and Music Computing, INESC TEC, Porto, Portugal
{gba, diogo.m.cocharro, carlos.guedes, mdavies}@inesctec.pt
[2] New York University Abu Dhabi, Abu Dhabi, United Arab Emirates

Abstract. We present Conchord, a system for real-time automatic generation of musical harmony through navigation in a novel 12-dimensional Tonal Interval Space. In this tonal space, angular and Euclidean distances among vectors representing multi-level pitch configurations equate with music theory principles, and vector norms acts as an indicator of consonance. Building upon these attributes, users can intuitively and dynamically define a collection of chords based on their relation to a tonal center (or key) and their consonance level. Furthermore, two algorithmic strategies grounded in principles from function and root-motion harmonic theories allow the generation of chord progressions characteristic of Western tonal music.

Keywords: Generative music · Harmony · Tonal pitch space

1 Introduction

Historically, Western tonal music has been subject to rigorous formalization, commonly expressed in a mathematical fashion. These formal tonal music frameworks are particularly adapted to computer modeling, which, in turn, can decisively contribute to machine musicianship [1]. A branch of this research is concerned with the automatic generation of harmony, which, within the Western tonal music context can be divided into two main problems: the automatic generation of chord progressions and the automatic harmonization of a given melody, of which the first concerns us here.

According to Wiggins [2], the automatic generation of musical harmony has been approached by different strategies, including (i) grammar-based systems [3], (ii) knowledge-based systems [4], (iii) biological-inspired algorithms [5, 6], (iv) constraint satisfaction systems [7] and (v) neural networks [8]. We extend Wiggins' taxonomy with a new category: statistical learning, which includes systems that generate harmony based on representations learned from musical examples [9, 10].

Central to this paper is the representation of the pitch configurations in generating musical harmony, especially through navigation in geometric pitch spaces, where neighborhood relations among the notes reflect perceptual or music theory properties.

R. Kronland-Martinet et al. (Eds.): CMMR 2015, LNCS 9617, pp. 243–260, 2016.
DOI: 10.1007/978-3-319-46282-0_15

The *Tonnetz* [11] is a planar representation of tonal pitch relations and one of the earliest examples of such geometrical spaces. Music theorists following the Riemann tradition adopted the Tonnetz to explain significant tonal pitch relationships, which are near one another in the space [12]. Chew's Spiral Array is another example of a geometric pitch space, which represents tonal pitch in a three-dimensional helix model and has been successfully applied to problems such as key estimation [13] and pitch spelling [14] from symbolic music data. Harte et al. [15] proposed the Tonal Centroid Space to estimate harmonic changes from musical audio by mapping 12-bin chroma vectors to the interior of a six-dimensional polytope. Recently, Bernardes et al. [16] presented a Tonal Interval Space, a twelve-dimensional space computed as the discrete Fourier transform of chroma vectors. While preserving the common-tone logic of the Tonnetz, the Tonal Interval Space overcomes an important limitation of the former by representing multi-level tonal pitch structures in the same space as unique vectors. Additionally, two important indicators can be computed from the space: tonal pitch relatedness and consonance.

Despite their intuitive and great explanatory potential, such spaces have yet to be widely applied to generative music. To the best of our knowledge, the only generative music systems built upon tonal pitch spaces were presented by: Behringer and Elliot [17] and Bigo et al. [18], who designed manually-driven generative systems based on the Tonnetz; Gatzsche et al. [19], who proposed a model for navigation on a three-dimensional space derived from Krumhansl and Kessler's [20] geometric representation of tonal pitch relations; and Chuan and Chew's [21] model for generating musical harmony through navigation in the Spiral Array.

Our approach to the automatic generation of musical harmony differs from, and extends previous research by proposing deterministic strategies for the real-time navigation in the Tonal Interval Space. We devise strategies to define spatial trajectories in the space for the generation of favorable chord progressions driven by principles from tonal harmony theories and build upon two main attributes of the space: tonal pitch relatedness and consonance. An important feature of our system is the possibility to modulate between keys—barely addressed in related software, yet a fundamental compositional strategy in tonal music.

Conchord, the prototype application of our model, was implemented for both Max [22] and Pure Data [23]. It allows expert and novice users to creatively explore musical harmony by assisting them in the automatic generation of 'generic' chord progressions characteristic of Western tonal music, rather than a specific response to a particular style or composer's idiom.

The main contributions beyond those in our previous work [24] include: (i) the use of a new set of weights in the Tonal Interval Space which better regulate the dimensions of the space for consonance measurements and (ii) the adoption of an algorithmic strategy to optimize the voice leading of selected chords given a user-defined pitch range.

The remainder of this paper is structured as follows. In Sect. 2 we present the architecture of Conchord, including the basic functioning of its component modules. In Sect. 3, we summarize the computation of the Tonal Interval Space in terms of its ability to measure tonal pitch relatedness among, and consonance level of, pitch configurations, in addition to the transposition invariance of tonal pitch directly mapped in the space. In Sect. 4, we describe the definition of a population of chords used during generation. In Sect. 5, we present the principles for optimal chord progression from two major harmony theories, which are then translated into spatial trajectories for the automatic generation of chord progressions in the Tonal Interval Space. In Sect. 6, we present an algorithmic strategy for unfolding selected chords represented by their component pitch classes to a user-defined pitch range while optimizing voice leading. In Sect. 7, we present the user interface of Conchord and the interaction design behind the application. In Sect. 8, we discuss the output of the system and finally, in Sect. 9, we state conclusions and future work.

2 Conchord: An Overview

Conchord is a system that automatically generates tonal harmony in real-time by navigation in a Tonal Interval Space, which places related pitch configurations at close distances. This property allows the definition of simple trajectories in the Tonal Interval Space driven by principles from function [25] and root-motion harmonic theories [26].

To enable navigation in the Tonal Interval Space and the generation of harmonic progressions, the system first populates the space with a collection of chords based on user input. These two main tasks are distributed into the following four modules of the system algorithmic chain: (i) initialization (ii) filtering, (iii) navigation/selection, and (iv) voice-leading (see Fig. 1).

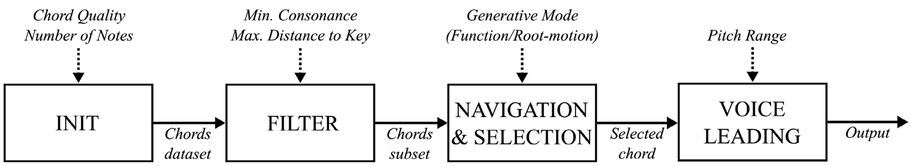

Fig. 1. Architecture of Conchord. The system modules are organized horizontally from left to right according to the information flow along with their input and output data.

The first module of the system automatically populates the Tonal Interval Space with a large set of chords based on user input specifying the chord quality and number of notes per chord. The second module is responsible for further constraining the chord dataset from the first module by filtering it according to two user-defined parameters: relatedness to a tonal center (or key) and consonance level. The resulting subset of chords is then stored into a database.

The third module comprises two generative algorithms for selecting favorable progressions among the (filtered) subset of chords through navigation in the Tonal Interval Space. The fourth and final module is responsible for organizing the voice leading of selected chords by unfolding its pitch class components into several pitch spacings and inversions over a user-defined pitch range. Then, a cost function ranks all generated solutions and outputs the candidate with the minimum cost for playback. Furthermore, this module allows the generation of chord progressions characteristic of the harmonic minor scale whenever generating chord progressions in a minor key and is responsible for formatting the internal pitch representation to MIDI note messages to be sent to, and synthesized by an external MIDI synthesizer.

3 Tonal Interval Space

The backend of Conchord relies on a 12-dimensional (12-D) Tonal Interval Space [16]. It represents multi-level pitch configurations (i.e., pitch classes, chords, and keys) on a geometric space where indicators of tonal pitch relatedness and consonance can be computed.

In this space, pitch configurations are represented by Tonal Interval Vectors (TIVs), *T(k)*, computed as the DFT of 12-D chroma vectors *c(n)* as follows:

$$T(k) = w(k)\sum_{n=0}^{N-1}\bar{c}(n)e^{-\frac{j2\pi kn}{N}}, \quad k \in \mathbb{Z} \quad \text{with } \bar{c}(n) = \frac{c(n)}{\sum_{n=0}^{N-1} c(n)}. \tag{1}$$

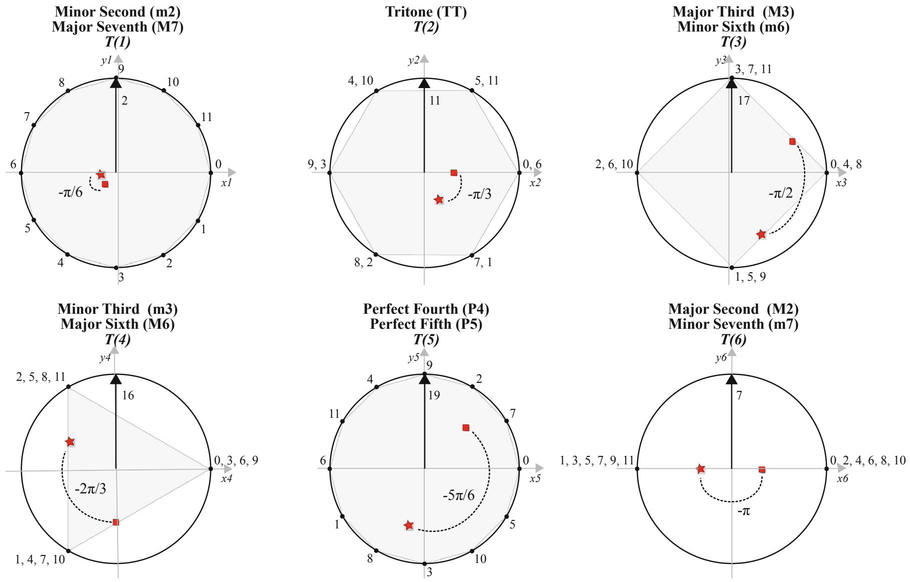

Fig. 2. Visualization of the Tonal Interval Space as six circles organized according to complementary intervals. Shaded grey areas denote the regions which TIVs can occupy for each circle. The plotted TIVs (square and star) correspond to the position of the C and C# major chords (pitch classes 0, 4, 7 and 1, 5, 8, respectively) in the space. The dashed circular lines indicate the angular rotation between C to C# major chords. The radii correspond to the weights of each complementary interval, *T(k)*, which for visualization purposes are all represented by identical size.

Here n is the chroma vector pitch class index up to $N = 12$, and $w(k) = \{2(\text{m2/M7}), 11(\text{TT}), 17(\text{M3/m6}), 16(\text{m3/M6}), 19(\text{P4/P5}), 7(\text{M2/m7})\}$ are weights derived from empirical ratings of dyads' consonance used to adjust the contribution of each dimension k of the space. $T(k)$ uses $\bar{c}(n)$, which is $c(n)$ normalized by the DC component $T(0) = \sum_{n=0}^{N-1} c(n)$ to allow the representation and comparison of all levels of tonal pitch represented by $c(n)$ in the same space. In the standard DFT calculation, we set $k = 12$ however, in practice, we only need to calculate T for coefficients 1 to 6, thus excluding the DC component and symmetrical coefficients.

Following Harte et al. [15], we visualize the 12-D space as 6 circles, each representing one complex DFT coefficient (see Fig. 2), i.e. circle 1 has the real part of *T(1)* on the x axis and the imaginary part of *T(1)* on the y axis and so on.

3.1 Chroma Vector Representation of Multi-level Tonal Pitch

For the purpose of this paper, we restrict musical notation to symbolic representations, and hence we express the presence of input pitch classes in the chroma vector $c(n)$ by binary activations. Each of the 12 elements of the chroma vector corresponds to a pitch class of the chromatic scale, assuming equal temperament and enharmonic equivalence.

For example, the C major chord (pitch classes 0, 4, and 7) gives the chroma vector $c = \{1, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0\}$. Major keys are represented by their diatonic set of pitch classes (e.g., C major results in the chroma vector $c = \{1, 0, 1, 0, 1, 1, 0, 1, 0, 1, 0, 1\}$), and minor keys by the diatonic set of their natural minor scale (e.g., C minor results in the chroma vector $c = \{1, 0, 1, 1, 0, 1, 0, 1, 1, 0, 1, 0\}$).

A particularity of our approach is the use of the natural minor scale, which shares the same pitch classes as its relative major scale, instead of the harmonic minor scale commonly adopted to represent the minor region. In doing so, we systematize the method behind the computation of chord sequences in the major and minor regions. Therefore, for the remainder of this paper, we will primarily refer to the major mode; moreover, in Sect. 6, we further detail a strategy for adopting the harmonic minor scale.

Please note that in the chroma vector $c(n)$, and consequently in the TIV $T(k)$, no information about pitch height is encoded. Thus, the octave cannot be represented nor suggested by $c(n)$ with binary encoding because all the octaves are collapsed into one. Additionally, the information relative to chord inversions and the effect of the relative amplitude between pitch classes are lost.

3.2 Spatial Relations Among Multi-level Tonal Pitch Configurations

The TIVs, $T(k)$, of different chroma vectors are separated by some spatial distance in the Tonal Interval Space. In this section, we discuss how these spatial distances can be understood within the most salient hierarchical layers of the tonal system from lower to higher levels of abstraction (i.e., pitch classes, then chords, and finally keys or regions), as well as how the levels interconnect.

Similar to the pitch organization of the Tonnetz [11], the Tonal Interval Space places pitch class intervals understood as related within the Western tonal music context at close distances, i.e., minimizing distances for unisons over perfect fifth over major/minor thirds, etc.

For this work, the most relevant feature of the Tonal Interval Space's chordal level is how it minimizes distances among chords sharing common tones. The greater the number of common tones between chords, the closer they are in the space, thus favoring chord progressions with minimal displacement of moving voices (known as voice-leading parsimony).

At the regional level, our space minimizes distances for closely related regions, where each key is flanked by its relative, dominant and subdominant keys. Additionally, the diatonic pitch class set of a given region is in the neighborhood of its key TIV. In other words, in-key pitch classes are at a smaller distance to their respective key TIV than outside-the-key pitch classes. Table 1 further demonstrates this property of the Tonal Interval Space by showing the angular and Euclidean distances between C major region TIVs and all 12 pitch classes.

Table 1. Angular (θ) and Euclidean (d) distances between C major regions TIVs and the entire set of pitch classes. The diatonic pitch class set of each region are in bold.

	C	C#	D	D#	E	F	F#	G	G#	A	A#	B
θ	**1.22**	2.12	**1.09**	2.12	**1.22**	**1.40**	2.02	**1.15**	1.97	**1.15**	2.01	**1.40**
d	**30.9**	39.63	**29.50**	39.63	**30.91**	**32.79**	30.80	**30.07**	38.41	**30.07**	38.80	**32.79**

3.3 Transposition

The geometry of our space allows us to explain and directly map transposition-related pitch configurations by rotation of TIVs. Transposing a pitch configuration by p semitones result in rotations of $T(k)$ by $\varphi(p) = \frac{-2\pi kp}{N}$ radians. In other words, the angle of rotation of TIVs is systematic and different for each of the complex DFT coefficients (i.e., each of the 6 circles shown in Fig. 2). For example, transposing a pitch configuration one semitone up is achieved by rotating the six coefficients of the TIV by −π/6, −π/3, −π/2, −2π/3, −5π/6, and −π, respectively.

Denoting $T_p(k)$ as the transposed version of $T(k)$ we have

$$T_p(k) = |T(k)|e^{-\frac{j2\pi(k+p)}{N}}. \tag{2}$$

Rotating a TIV through all 12 semitones of the chromatic scale creates a concentric 'layer' containing all of its 12 possible transpositions. Therefore, pitch configurations sharing the same interval vector (e.g., chords with the same quality) are at the same distance from the center.

3.4 Measuring Consonance

Two important design principles of the Tonal Interval Space allow the computation of a tonal pitch consonance indicator from the space. They are the normalization by $T(0)$

and the weights $w(k)$ applied in Eq. (1). The former constrains the space to a limited area where all multi-level pitch configurations a chroma vector can represent exist. In the proposed visualization of the 12-D space shown in Fig. 2, this is equivalent to saying that all pitch configurations lie inside the 6 circles. The latter distorts the DFT coefficients to regulate the contribution of each interval according to empirical ratings of consonance. These two elements combined create a space in which pitch classes at the edge of the space and furthest from the center are considered the most consonant configurations, and a configuration with all 12 active pitch classes in the center of space is considered the most dissonant. These explicit design principles orders the norm of TIV, $\|T(K)\|$, for both dyads (i.e., P1 > P4/P5 > m3/M6 > M3/m6 > TT > M2/m7 > m2/M7) and common triads (i.e., major/minor > suspended fourth > diminished > augmented) according to empirical data derived from previous listening experiments [27–31].

Hence, we extrapolate the consonance measure *C(T)* of the $\|T(K)\|$ by the norm $T(k)$, which can also be calculated as the Euclidean distance from the centre as follows:

$$C(T) = \|T(K)\| = \sqrt{T(k) \cdot T(k)} = \sqrt{\sum_{k=1}^{M} |T(k)|^2}. \quad (3)$$

4 Defining a Population of Chords

The first module of our algorithmic chain defines a population of chords to be used during generation by a twofold strategy. First, the system populates the space with a large set of chords with variable numbers of notes and qualities. Second, those chords that correspond to a pre-defined tonal center and minimum consonance level are retained and all others are discarded.

Conchord includes four options to initially populate the Tonal Interval Space with chords: (i) all *m*-note chords combinations without repetition from the set of 12 pitch classes; (ii) all major, minor, and augmented triads; (iii) all major, minor and half-diminished seventh chords; and (iv) the manual specification of chords.

After creating an initial population of chords, T_i, these are reduced to a subset of chords, Z_i, by applying user-defined constraints that regulate the chords' relation to a key and the chords' consonance. A maximal distance of the chords from a key TIV defines the key-relatedness constraint, and a minimum norm of the chords TIV defines the consonance constraint. In this way, the subset of chords, Z_i, meets both conditions simultaneously:

$$Z_i = T_i\{D_i < D_{max} \,\&\, C(T_i) > C_{min}\}, \quad (4)$$

where,

$$D_i = \sqrt{\sum_k [|T_{key}(k) - T_{chord}(k)|]^2}, \quad (5)$$

D_{max} is the maximum distance of chord TIVs, $T_{chord}(k)$, from a key TIV, $T_{key}(k)$, and C_{min} denotes the user-defined minimal chord consonance. Both constraints can be assigned using Conchord's interface as described in section. Typical ranges for these parameters are $11 < D_{max} < 33$ and $15 < C_{min} < 25$. The TIVs and the component pitch classes of the resulting subset of chords are then stored in a database.

The modulation between keys in the Tonal Interval Space combines the possibility to transpose keys by TIV rotation with the constraint-based strategies stated in (5). Given both the interval distance between a target and departure keys in semitones and their TIVs, we can modulate between them by defining an angular trajectory in the Tonal Interval Space that continuously discards and adds new chords to the chord database. This angular trajectory corresponds to the interpolation between target and departure key TIVs given by the smaller angular rotation (either clockwise or anti-clockwise) on each individual circle calculated using (3). The resulting trajectory enssures that overlapping and adjacent pitches and chords across regions favor a smooth transition between keys.

5 Chord Candidate Selection

This section details two algorithmic strategies for defining 'optimal' motions in the Tonal Interval Space among the database chords, Z_i. Chord selection occurs once every beat and is triggered by an internal clock adjusted to a user-defined tempo.

The algorithmic strategies for navigation in the Tonal Interval Space are driven by formalized principles from two theories of tonal harmony: function and root-motion harmonic theories (Sect. 5.1). The translation of formalized principles to trajectories in the Tonal Interval Space for the generation of favorable chord progressions is detailed in Sects. 5.2 and 5.3.

5.1 Theories of Tonal Harmony

Tonal harmony has been a frequently visited topic by music theorists since Rameau's original treatise [32]. Among the most widely accepted and popular theories of harmony we find: (i) scale-degree theories, which claim that each chord can resolve in its own characteristic way; (ii) function theories, which group chords into larger categories; and (iii) root-motion theories, which emphasize the intervals formed between successive chord-roots.

The principles for favorable chord progressions stated by the two latter theories will be presented next, with the aim of inferring formalisms which can then be translated into trajectories in the Tonal Interval Space. We do not address scale-degree theories because their underlying principles cannot be modeled in our space due to their asymmetric nature. For a comprehensive explanation of these theories please refer to [33].

Function Theories. Originated by Riemann [25], function theories examine and explain musical harmony according to two fundamental components: chord categories based on shared harmonic function and typical motions between these categories [34].

In detail, the first component groups chords together into three categories: (i) 'tonic,' (ii) 'subdominant,' and (iii) 'dominant' [25]. The tonic category comprises the tonic, mediant, and submediant degrees (the ***I***, ***iii***, and ***vi*** degrees); the subdominant category comprises the subdominant, supertonic, and submediant degrees (the ***IV***, ***ii***, and ***vi*** degrees); and the dominant category comprises the dominant, leading-tone, and mediant degrees (the ***V***, ***vii***°, and ***iii*** degrees). The second component of the function theory postulates normative patterns of functional progressions. According to Riemann [25], usual motions depart from the tonic to the subdominant to the dominant and then back to the tonic.

Root-Motion Theories. Following the seminal work of Rameau [32], root-motion theories have been extended by Schoenberg [26] and Meeus [35]. Root-motion theories claim that standard and common tonal progressions can be defined in terms of the intervals formed by the root of consecutive chords, emphasizing a set of 'optimal' or common chord root progressions, compared to 'atypical' sequences.

The guidelines for optimal chord progressions in terms of the type of root motion adopted here are based on Schoenberg [26], who distinguishes three root progression cases: 'ascending,' 'descending' and 'super-strong' progressions. In an ascending progression, the chord root moves a fourth up or a third down (e.g., ***I–IV*** and ***I–vi***, respectively). In a descending progression the chord root moves up a fifth or a third (e.g., ***I–V*** and ***I–iii***, respectively). Finally, in a super-strong progression a chord root moves a second up or down (e.g., ***V–vi*** and ***V–IV***, respectively).

While encouraging the unreserved use of ascending root progressions, Schoenberg discourages the use of descending progressions—which should be treated as passing chords between ascending progressions—and the sparse use of super-strong progressions.

5.2 Angular Trajectories Driven by Harmonic Function Theory Principles

To define angular trajectories in the Tonal Interval Space that convey the harmonic function theory principles stated above, we first define three harmonic function regions—tonic, subdominant, and dominant—in the 12-D space. These harmonic function regions are represented by three vectors departing from the key TIV, $T_{key}(k)$, to the tonic, subdominant, and dominant TIVs, $T_i(k)$, respectively (represented in Fig. 3 as $\vec{T}_{\mathbf{I}}$, $\vec{T}_{\mathbf{IV}}$, and $\vec{T}_{\mathbf{V}}$), such that

$$\vec{T}_i(k) = T_i(k) - T_{key}(k). \tag{6}$$

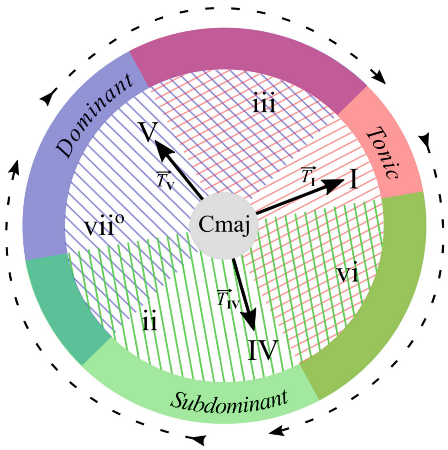

Fig. 3. 2-D visualisation of the 12-D representation of the C major key and its common component diatonic triads using nonmetric multidimensional scaling—Roman numerals label the harmonic function of each triad. The pattern regions denote three harmonic function categories (i.e., Tonic, Subdominant, and Dominant) and $\vec{T}_{\mathbf{I}}$, $\vec{T}_{\mathbf{IV}}$, and $\vec{T}_{\mathbf{V}}$ their representative vectors. Dashed circular lines represent typical motions between harmonic function categories (i.e., Tonic to Subdominant to Dominant and back to Tonic).

For the purpose of visualization of the harmonic function categories regions we use a 2-D representation of the 12-D Tonal Interval space obtained by nonmetric multidimensional scaling[1] (MDS) as shown in Fig. 3.

Relying on the vector representations for the three harmonic function regions, we then define in the 12-D space favorable motions between them following Riemann [25]. To this end, we sequentially alternate between the target vectors $\vec{T}_{target}$ (i.e., $\vec{T}_{\mathbf{I}}$, $\vec{T}_{\mathbf{IV}}$, and $\vec{T}_{\mathbf{V}}$) and randomly select a chord from the database within a $\pi/2$ angular distance, computed by the angle ϑ_i in Eq. (7). To perform this operation, all database chord candidates, $T_i(k)$, are first converted to vectors departing from the key TIV using Eq. (6).

$$\vartheta_i = \cos^{-1} \frac{\vec{T}_{target} \cdot \vec{T}_i}{\| \vec{T}_{target} \| \cdot \| \vec{T}_{i} \|}. \tag{7}$$

[1] Sheppard [36] and Kruskal [37] first used this method, which has been extensively applied to visualise representations of multidimensional pitch structures [38, 39]. Briefly, nonmetric MDS attempts to transform a set of n-dimensional vectors, expressed by their distance in the item-item matrix, into a spatial representation that exposes the interrelationships among a set of input cases. We use the `smacof` library [40] from the statistical analysis package 'R' to compute dimensionality reduction using a nonmetric MDS algorithm. More specifically, we use the function `smacofSym`, with 'ordinal' type and 'primary' ties.

5.3 Angular Trajectories Driven by Root-Motion Harmonic Theory Principles

We define root-motion motion between successive chords in a stochastic manner based on angular distances. Prior to selection, the system calculates the vectors departing from an assigned key TIV, $T_{key}(k)$, to the chord database TIVs, $T_i(k)$, and to an additional set of seven diatonic triads TIVs for each scale degree—referred hereafter as the target vector, $\vec{T}_{\mathbf{I}} - \vec{T}_{\mathbf{vii}^\circ}$, where the chord database and target vectors are computed using (6).

At each new chord selection, a stochastic algorithm decides the next chord root following Schoenberg's suggestions for favorable chord-root progressions, i.e. favoring ascending progressions, over super-strong progressions, over descending progressions—according to heuristically assigned probabilities of 60 % (ascending), 20 % (super-strong), 10 % (descending) and 10 % (remaining in the same chord root).

Figure 4 shows the organization of the diatonic triad TIVs and their representative (target) vectors around the key TIV, as well as the probabilities of transition between them departing from the root of the tonic (left) and subdominant (right) degrees.

Then, given a chord root for the next progressions, the system identifies the corresponding triad with such a root and retrieves its representative target vector. Finally, the system computes in the 12-D space the angle ϑ_i between the target vector and the chord database vectors using (7), and randomly selects a chord from the set of database chord within the range $\pi/6$ around the target angle.

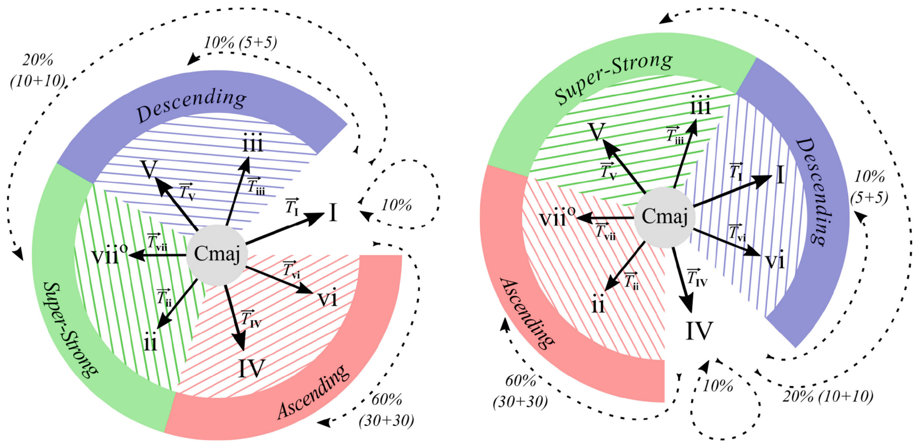

Fig. 4. 2-D visualisation of the 12-D representation of the C major key and its common component diatonic triads using nonmetric multidimensional scaling—Roman numerals label the harmonic function of each triad. The pattern regions organize and highlight the three root-motion progressions proposed by Schoenberg [26], departing from the tonic (left) and subdominant (right). Bold lines correspond to the vectors $\vec{T}_{\mathbf{I}} - \vec{T}_{\mathbf{vii}^\circ}$ of common diatonic triads from their key TIV and dashed lines the probability of transitioning between root-motion cases.

6 Voice Leading and Playback

Voice leading is the term used to denote the horizontal motion of the individual parts (or voices) in harmony writing. The rules for voice leading together with the rules for harmonic progression establish the ground for successful musical harmony writing. Its importance is highlighted by Schonbrun [41, p. 174], who claims that 'good voice leading can take a simple chord sequence and transform it into a masterpiece'.

A vast amount of literature detailing objective rules for optimal voice leading exists [42]. The most common rules for voice leading found in music theory textbooks are drawn from, and meant for, four-part vocal music. Even though these rules can be applied to music outside this category, we should bear in mind that voice leading is highly linked and commonly adapted to the context of the piece, including style, orchestration, composer idiosyncrasies, etc.

We follow a strategy first detailed in [16] to enhance the voice leading of selected chords given their pitch class sets and a user-defined pitch range. The rules adopted here were selected from a large collection of voice leading rules framed for four-part vocal music in [42]. The selection criteria were based on the adaptability of the rules to general tonal music contexts other than vocal music, resulting in the following five rules:

1. Vertical (chord) spacings—intervals between parts should not exceed one octave, with the exception of the interval formed by the two lowest parts, in which no restriction is applied;
2. Avoid large horizontal (or melodic) leaps—adjacent chords should minimize the intervallic distance in each voice;
3. Avoid parallel octaves and fifths—intervals of the octave and fifth should not happen consecutively in the same parts;
4. Avoid hidden octaves and fifths—intervals of the octave and fifth should not happen consecutively in different parts;
5. Outer parts contrary motion—contrary motion between outer parts is encouraged.

Based on the above-stated rules, we devised an algorithmic strategy that finds the best voice leading for selected chords from a large set of chord candidates generated from unpacking the set of pitch classes to all possible inversions and vertical spacings within a user-defined pitch range. For example, a 3-note chord will result in a set of chord candidates which include the fundamental state as well as its two inversions (i.e. with the third and fifth on the bass). Additionally, for each of the three sets all possible chord spacings within the user-defined pitch range are instantiated.

All candidate chords generated are evaluated in relation to the previously played chord and are ranked using the set of points shown in Table 2. An algorithm assigns a cost for each of the five voice leading rules per chord candidate. Then, an overall cost per chord is computed by summing the five rule condition costs. Finally, the algorithm outputs the chord with the minimum (i.e., best) voice leading cost.

After a chord has been selected it is converted to MIDI note messages for playback using a synthesizer. For simplicity, we assign an arbitrary MIDI note velocity of 100 for all notes and set the duration to be equal to the length of each beat.

Table 2. Point assignments for voice leading rule conditions.

		1. Chord spacing	2. Melodic leap	3. Contrary motion	4. Parallel 5th/8th	6. Hidden 5th/8th
Cost	5	–	–	–	true	true
	2	>12	>8	–	–	–
	1	>3	>4	false	–	–
	0	=0	=0	true	false	false

Prior to the definition of the chord candidate that best fits the voice leading rules while in a minor key, a parallel operation shifts the seventh degree of the natural minor scale a semitone higher (e.g. on a minor scale, shifting the G natural to G sharp). While the use of the natural minor scale simplifies the generations and database construction, this additional processing allows the generation of chords characteristic of the harmonic minor scale, the most common scale for chord construction in minor keys.

7 User Interface

To allow users to interact with Conchord, we have created an interface for both Max [22] and Pure Data [23], of which a screenshot of the latter is shown in Fig. 5. Conchord's interface was designed to provide a simple, intuitive, and flexible experience for both expert and novice users. Despite minor cosmetic differences, the Max and Pure Data versions have identical functionality.

The interface is composed of three modules stacked vertically. The two upper modules are responsible for (i) setting the parameters for initializing and filtering out

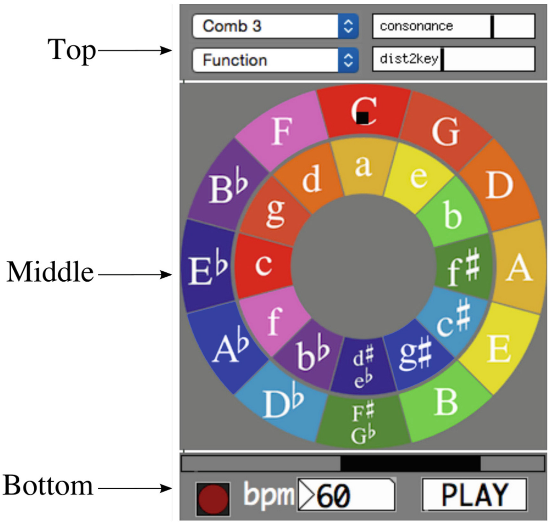

Fig. 5. Conchord interface developed for the PD programming environment.

the population of chords as detailed in Sect. 4 and for (ii) defining the generative strategy behind chord selection between function and root-motion harmonic theory principles. The lower module is responsible for controlling the playback as well as specifying the pitch range chords can occupy.

To initialize the Tonal Interval Space with a population of chords, the user must first define on the top-left menu of the uppermost module the type of chord aggregates initially generated by selecting one of following four predefined options: (i) 'Diatonic triads,' (ii) 'Diatonic tetrads,' (iii) 'Comb 3'—which stand for all 3-note chords combinations without repetition from the 12 pitch classes, and (iv) 'Comb 4'—which stands for all 4-note chords combinations without repetition from the 12 pitch classes. Then, the user should specify in the middle module a key among the 24 minor and major keys represented on a double circle of fifths, which mirrors the key organization of the Tonal Interval Space. The selected key will be then used to calculate the chord relatedness to a tonal center.

The first two top-left menu entries automatically initialize the Tonal Interval Space with the common diatonic triads and tetrads of a given key by populating the space with all major, minor and diminished chords and further filtering them according to pre-defined key-relatedness and consonance parameter values. Top-left menu entries three and four offer a greater degree of freedom and allow users to initialize the space with all combinations of 3- and 4-note chords, as well as the possibility to manually specify both filtering parameters. Furthermore, the space can also be populated by a set of chords manually chosen via a standard Max and Pure Data message format specifying their pitch class components. In the bottom-left menu of the top interface module, the user can define one of two 'Function and 'Root-motion' generative modes.

Finally, the bottom module of the interface includes parameters that regulate the pitch range chords can occupy, the tempo (in beats-per-minute), as well as start/stop and recording commands. The output format of the recoding is in the standard MIDI format, which can be further edited in traditional MIDI sequencers.

8 Results

This section discusses musical results generated by the two strategies implemented in Conchord, of which a transcription along with its function and root-motion harmonic analysis is shown in Fig. 6.

Figure 6 includes two staves, resulting from each generative strategy: the upper one corresponding to the harmonic function theory and the lower to the root-motion harmonic theory. Additionally, each stave is split into two parts, which differ according to the chord database. The first part only adopts common (major, minor and diminished) diatonic triads of the C major key. The second part remains in the key of C major, but adopts an extended degree of dissonance in relation to the core diatonic triads.

In the particular case of the harmonic function theory (upper stave of Fig. 6) the chord sequence follows the functional category motions proposed by Riemann [25], i.e. the normative patterns of functional progressions from the tonic (T) to the subdominant (SD) to the dominant (D) and back to the tonic. Additionally, the sequences of functional harmonic categories are not disturbed in the second part of the stave when

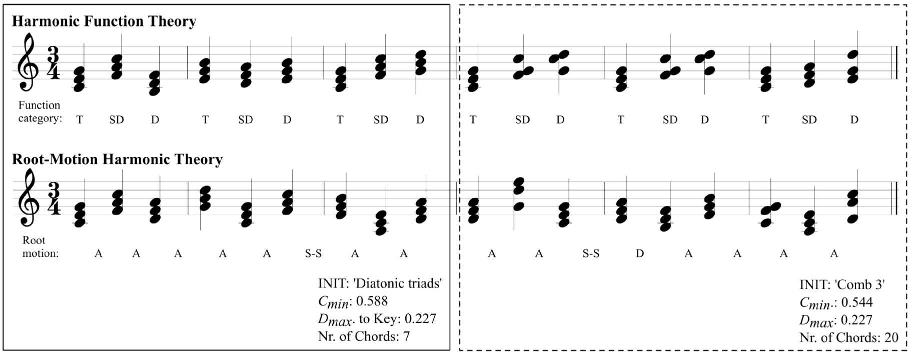

Fig. 6. Chord progressions generated with Conchord. The upper stave was generated using the harmonic function theory principles and includes an analysis of the chord's functional category (i.e., Tonic, SubDominant, and Dominant). The lower stave was generated according to root-motion theory principles and includes an analysis of the root-motion cases (i.e., Ascending, Descending, and Super-Strong). The bold square circumscribes the use of common diatonic triads of the C major key, and the dashed square remain in the key of C major but extends the database while adopting an extended degree of dissonance. To facilitate reading, chords were rearranged to their root position.

decreasing the level of consonance (which results in a higher degree of tonal ambiguity).

In the lower stave of Fig. 6, the chord progressions obey Schoenberg's [26] guidelines for favorable chord-root motions. Note that in the generated excerpt, ascending progressions occur 76 % of the time, super-strong progressions were explicitly restricted to 12 % and descending progressions or remaining in the same chord root to 6 % each. As in the function harmonic theory, decreasing the level of dissonance didn't disturb the expected the probability of occurrence of chord-root motions.

To summarize, the generated chord sequences demonstrate the efficiency of the algorithms, which convey the principles of the function and root-motion harmonic theories, as well as the behavior of the generative music algorithms under more dissonant and less representative chord aggregates of a tonal center. Several more musical examples generated by the system along with supplementary harmonic analysis and expanded use of generated chord progressions via arpeggiation and orchestration are available at: http://smc.inesctec.pt/technologies/conchord.

9 Conclusions and Future Work

In this paper we presented strategies for navigating in the Tonal Interval Space, a novel 12-D representation for the automatic generation of chord progressions based on harmonic theory principles. The primary contributions of our Tonal Interval Space in comparison to existing methods, notably those by Chew [13] and Harte et al. [15], are the possibility to explain the relations among pitch configurations at the three fundamental pitch, chord and key levels of Western tonal music in a single space, and to provide an indicator of the consonance of pitch configurations.

The algorithm and applications presented establish a novel approach to the real-time automatic generation of harmony driven by the tonal pitch organization of the Tonal Interval Space in combination with principles derived from function and root-motion harmonic theories. The algorithms convey the principles established by both theories, which have been modeled in existing generative music applications. The novelty of our approach relies on the highly flexible and dynamic specification of a chord database prior to generation based on user-defined chord parameters including: the number of notes, chord quality, relatedness to a key, and level of consonance, which in itself expands upon existing systems for the automatic generation of chord progression beyond the common four-voice setting and conventional use of simple triads.

While we only formalized principles for favorable chord progressions in this paper, many musical aspects such as rhythmic structure are not modeled currently. A more refined model will be addressed in future work by incorporating additional musical components, including the integration of the current algorithmic strategies within a metrical template. By doing so, we will also be able to more efficiently define strategies for selecting chords within the large chord categories considered by our system.

Acknowledgments. This research is financed by National Funds through the FCT - Fundação para a Ciência e a Tecnologia within post-doctoral grants SFRH/BPD/109457/2015 and SFRH/BPD/88722/2012.

References

1. Rowe, R.: Machine Musicianship. MIT Press, Cambridge (2001)
2. Wiggins, G.A.: Automated generation of musical harmony: what's missing. In: Proceedings of the International Joint Conference in Artificial Intelligence (1999)
3. Rohrmeier, M.: A generative grammar approach to diatonic harmonic structure. In: Proceedings of the 4th Sound and Music Computing Conference, pp. 97–100 (2007)
4. Pachet, F.: The MusES system: an environment for experimenting with knowledge representation techniques in tonal harmony. In: 1st Brazilian Symposium on Computer Music, pp. 195–201 (1994)
5. Phon-Amnuaisuk, S., Tuson, A., Wiggins, G.: Evolving musical harmonisation. In: Dobnikar, A., Steele, N.C., Pearson, D.W., Albrecht, R.F. (eds.) Artificial Neural Nets and Genetic Algorithms, pp. 229–234. Springer, Vienna (1999)
6. Navarro, M., Caetano, M., Bernardes, G., de Castro, L.N., Corchado, J.M.: Automatic generation of chord progressions with an artificial immune system. In: Johnson, C., Carballal, A., Correia, J. (eds.) EvoMUSART 2015. LNCS, vol. 9027, pp. 175–186. Springer, Heidelberg (2015)
7. Pachet, F., Roy, P.: Formulating constraint satisfaction problems on part-whole relations: the case of automatic musical harmonization. In: ECAI 98 Workshop on Constraints for Artistic Applications, pp. 1–11 (1998)

8. Gang, D., Lehmann, D., Wagner, N.: Harmonizing melodies in real-time: the connectionist approach. In: Proceedings of the International Computer Music Association, Thessaloniki, Greece, pp. 27–31 (1997)
9. Manaris, B., Johnson, D., Vassilandonakis, Y.: Harmonic navigator: a gesture-driven, corpus-based approach to music analysis, composition, and performance. In: Proceedings of the 9th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, pp. 67–74 (2013)
10. Eigenfeldt, A., Pasquier, P.: Realtime generation of harmonic progressions using controlled markov selection. In: Proceedings of ICCC-X-Computational Creativity Conference, pp. 16–25 (2010)
11. Cohn, R.: Neo-riemannian operations, parsimonious trichords, and their "tonnetz" representations. J. Music Theory **41**(1), 1–66 (1997)
12. Cohn, R.: Introduction to neo-riemannian theory: a survey and a historical perspective. J. Music Theory **42**(2), 167–180 (1998)
13. Chew, E.: Towards a mathematical model of tonality. Ph.D. dissertation, MIT (2000)
14. Chew, E., Chen, Y.: Determining context-defining windows: pitch spelling using the spiral array. In: Proceedings of the International Society for Music Information Retrieval Conference (2003)
15. Harte, C., Sandler, M., Gasser, M.: Detecting harmonic change in musical audio. In: Proceedings of the 1st ACM Workshop on Audio and Music Computing Multimedia, pp. 21–26. ACM, New York (2006)
16. Bernardes, G., Cocharro, D., Guedes, C., Davies, M.E.P.: Harmony generation driven by a perceptually motivated tonal interval space. In: ACM Computers in Entertainment. ACM, New York (2015, in press)
17. Behringer, R., Elliot, J.: Linking physical space with the Riemann Tonnetz for exploration of western tonality. In: Hermida, J., Ferrero, M. (eds.) Music Education. Nova Science Publishers Inc, Hauppauge (2009)
18. Bigo, L., Garcia, J., Spicher, A., Mackay, W.E.: PaperTonnetz: music composition with interactive paper. In: Proceedings of the 9th Sound and Music Computing Conference, pp. 219–225 (2012)
19. Gatzsche, G., Mehnert, M., Stöcklmeier, C.: Interaction with tonal pitch spaces. In: 8th International Conference on New Interfaces for Musical Expression, Genova, Italy, pp. 325–330 (2008)
20. Krumhansl, C.L., Kessler, E.J.: Tracing the dynamic changes in perceived tonal organization in a spatial representation of musical keys. Psychol. Rev. **89**, 334–368 (1982)
21. Chuan, C.-H., Chew, E.: A hybrid system for automatic generation of style-specific accompaniment. In: Proceedings of the 4th International Joint Workshop on Computational Creativity, pp. 57–64. Goldsmiths, University of London (2007)
22. Pure Data. http://puredata.info
23. Max. https://cycling74.com
24. Bernardes, G., Cocharro, D., Guedes, C., Davies, M.E.P.: Conchord: an application for generating musical harmony by navigating in a perceptually motivated tonal interval space. In: Proceedings of the 11th International Symposium on Computer Music Modeling and Retrieval (CMMR), pp. 71–86 (2015)
25. Riemann, H.: Vereinfachte Harmonielehre. Augener, London (1893)
26. Schoenberg, A.: Structural Functions of Harmony, 2nd edn. W. W. Norton Inc., New York (1969). [1954], (Revised by Stein, L.)
27. Malmberg, C.F.: The Perception of consonance and dissonance. Psychol. Monogr. **25**(2), 93–133 (1918)

28. Kameoka, A., Kuriyagawa, M.: Consonance theory. Part I: consonance of dyads. J. Acoust. Soc. Am. **45**, 1451–1459 (1969)
29. Hutchinson, W., Knopoff, L.: The acoustic component of Western consonance. Interface **10**(2), 129–149 (1979)
30. Roberts, L.A.: Consonance judgements of musical chords by musicians and untrained listeners. Acta Acustica United Acustica **62**(2), 163–171 (1986)
31. Cook, N.D.: Harmony, Perspective, and Triadic Cognition. Cambridge University Press, New York (2012)
32. Rameau, J.-P.: Traité de l'Harmonie. Ballard, Paris (1722). Treatise on Harmony (trans: Gossett, P.) Dover, New York (1971)
33. Tymoczko, D.: Progressions Fondamentales, Fonctions, Degrés: Une Grammaire de l'Harmonie Tonale Élémentaire. Musurgia: Analyse et Pratique Musicales **10**(3–4), 35–64 (2003)
34. Agmon, E.: Functional harmony revisited: a prototype-theoretic approach. Music Theory Spectr. **17**(2), 196–214 (1995)
35. Meeus, N.: Toward a post-Schoenbergian grammar of tonal and pre-tonal harmonic progressions. Music Theory Online **6**(1) (2000)
36. Shepard, R.N.: Structural representations of musical pitch. In: Deutsch, D. (ed.) The Psychology of Music, pp. 335–353. Academic Press, New York (1982)
37. Kruskal, J.B.: Nonmetric multidimensional scaling: a numerical method. Psychometrika **29**, 28–42 (1964)
38. Krumhansl, C.L., Kessler, E.J.: Tracing the dynamic changes in perceived tonal organization in a spatial representation of musical keys. Psychol. Rev. **89**, 334–368 (1982)
39. Lerdahl, F.: Tonal Pitch Space. Oxford University Press, New York (2001)
40. De Leeuw, J., Mair, P.: Multidimensional scaling using majorization: SMACOF in R. J. Stat. Softw. **31**(3), 1–30 (2009)
41. Schonbrun, M.: The Everything Music Theory: Take Your Understanding of Music to the Next Level. Adams Media, Avon (2011)
42. Huron, D.: Tone and voice: a derivation of the rules of voice-leading from perceptual principles. Music Percept. **19**(1), 1–64 (2001)