

Thinking Out of the Grid and Inside the Spiral – Geometric Interpretations of and Comparisons with the Spiral Array Model

ELAINE CHEW

EPSTEIN DEPARTMENT OF INDUSTRIAL AND SYSTEMS ENGINEERING
HSIEH DEPARTMENT OF ELECTRICAL ENGINEERING
VITERBI SCHOOL OF ENGINEERING
UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CALIFORNIA, USA

RADCLIFFE INSTITUTE FOR ADVANCED STUDY (2007-2008 FELLOW)
CAMBRIDGE, MASSACHUSETTS, USA

Abstract

This paper presents a geometric interpretation of the Spiral Array model and its comparison to Lerdahl's Tonal Pitch Space and Krumhansl's spatial representation of pitch relations. The Spiral Array model is derived from a three-dimensional configuration of the Harmonic Network. The fundamental idea underlying the model is the representing of higher-level objects in the spiral's interior as convex combinations of the representations of the lower level components. By using the interior of the spiral, the original discrete space is relaxed to one that is continuous. Geometric mappings are demonstrated among Lerdahl's Tonal Pitch Space, Krumhansl's and Krumhansl & Kessler's spatial representations of pitch class and key relations, and the Spiral Array model. The interior-point approach of the Spiral Array model is shown to generate higher-level structures that are consistent with the results of these other approaches. The advantages of the interior-point approach are that it facilitates comparisons across different hierarchical levels, and problems that were previously combinatorial in nature can be modeled more efficiently and robustly, mathematically and computationally, using the continuous space in the interior.

Author's estimated word count (<5,000): 6090

Date of submission: 15 October, 2003

Date of revision: 12 September, 2007

1 Introduction

Tonality is the system of relationships that generates a hierarchy among pitches, resulting in one pitch being the most stable. Numerous geometric models for these pitch relations have been proposed. According to Shepard (1982a), any “cognitive representation of musical pitch must have properties of great regularity, symmetry, and transformational invariance.” Not surprisingly, many of these models are based on lattices that can be wrapped around cylinders to form helical structures. One such lattice is the Harmonic Network, also known as the *tonnetz*. The Harmonic Network clusters pitch classes that form higher-level structures in the tonal system, such as triads and keys. The lattice repeats periodically in such a way that one can roll it, at an angle, onto a cylinder so that the repeating pitch class names line up one over another.

This paper argues for the representing of higher-level objects, in a systematic fashion, as spatial points inside such a pitch class cylinder generated by the Harmonic Network. The continuous three-dimensional space inside the spiral provides a metric for quantifying the distance between any two objects represented in the same space. Consider the pitch class representations, the vertices of this network, and the edges in the network that connect these vertices. These edges between the vertices mark the distance between any two pitch classes represented on the network.

In its original form, the Harmonic Network cannot easily provide a consistent metric for measuring the distance between objects that generalizes to higher-level entities beyond pitch classes. By venturing inside the cylinder, one can represent triads and keys in the same three-dimensional space as that for the pitch classes, and measure the distance between any two objects, even those from different hierarchical levels, thus, in a sense, treating all objects equally. The idea of representing higher-level objects inside the pitch class cylinder generalizes so that not only traditional objects such as triads and keys, but also less traditional pitch sets, can be represented and compared quantitatively in the interior space.

As mentioned above, a distinct advantage of the interior approach is that objects from different hierarchical levels are represented in the same space, thus facilitating inter-level object comparisons. By utilizing the interior space, problems of pattern recognition, such as chord recognition and tonal induction, can be reduced from a combinatorial one to a simple nearest neighbor search, as in the cases of key-finding (Chew 2001) and pitch spelling (Chew & Chen 2005). Furthermore, the interior-point

approach provides a distance metric that allows for the design of computationally efficient algorithms for problems of tonal comparison and change detection (see Chew 2006, and the paper by Volk & Chew in this issue.)

The inspiration for this interior-point approach came from the field of Operations Research. In the domain of linear optimization, for many decades, the method of choice to solve linear programming problems was the Simplex Method, invented in 1947 by George Dantzig (1963). A linear optimization problem is one of finding the optimal solution as measured by maximizing or minimizing a linear objective function, while satisfying a set of linear constraints. The constraints can be viewed as hyper-planes in higher-dimensional space that form the borders of a feasible region or solution space. The optimal solution, when one exists, resides at one of the corner points of the feasible region defined by these constraints. The Simplex Method finds the optimal solution by pivoting between adjacent corner point solutions, through the edge that gives the fastest rate of improvement in the objective function. The pivoting between adjacent corner point solutions is akin to transitions between neighboring chords in the dual graph of the Harmonic Network. The Harmonic Network and its dual graph will be described in the next section.

Even though the Simplex Method has a computational complexity that is exponential, it proves to be a reasonable approach in practice. In 1984, Karmarkar proposed the interior-point approach (see Hillier & Lieberman 2001, p.163–168). In interior-point approaches, rather than pivoting on the vertices of the feasible region in search of optimal solutions, the corner point requirement is relaxed to allow for the search to proceed in the interior space, even though the optimal solution will necessarily be at a corner point. When iterations of the process lead to only minute improvements in the objective function, the closest corner point solution is selected as the optimal solution. By traveling through the interior of the solution space, rather than pivoting through plausible solutions, albeit in a smart way, interior-point algorithms have been shown to be polynomial in complexity.

Likewise, the Spiral Array approach relaxes the requirement to stay on the vertices of the Harmonic Network, or those of its dual chord space, to seek best solutions in the interior space. In the spirit of the adage "a picture's worth a thousand words," the first part of the paper presents an image-driven guide to the geometry of the Spiral Array model, first proposed in (Chew 2000). Representing objects out of the grid and inside the spiral is the fundamental idea behind the Spiral Array model. The second part of the paper compares the Spiral Array's geometric

structures with other spatial representations of tonal pitch space proposed by Krumhansl (1978, 1982, 1990) and Lerdahl (2001), and discusses the similarities and differences between the metrics employed. Mappings among the three models are demonstrated geometrically, validating the use of the interior-point approach to representing tonal objects in space.

2 From Grid to Spiral Representation

The Harmonic Network is a network representation of pitch relations where each node represents a pitch class, that is to say, a set of pitches related by some multiple of an octave. In network terminology, each node is a vertex of degree six (having six edges incident on the node). Each opposing pair of edges connects the pitch class to other nodes related by one of three intervals – Perfect fifth (P5), major third (M3) and minor third (m3) – as shown in Figure 1. The Harmonic Network, also known as the *tonnetz*, forms the foundation of Neo-Riemannian Theory and has been attributed to the mathematician Euler (see Cohn 1998a, Lewin 1982 and 1987).




Figure 1. The Harmonic Network, also known as the *Tonnetz*.

Each triangle in the Harmonic Network forms a triad (major or minor depending on its orientation). The network of triads forms the dual graph of the Harmonic Network (see Figure 2). Each new edge that cuts across an arc in the original lattice represents a distance-minimizing transformation between two triads, a transformation that exhibits the property of parsimonious voice leading. Transformations on the dual graph have been used to analyze triadic movement in tonal and post-tonal music (see Cohn 1996 and 1997).

Ftonnetz2-cim.eps

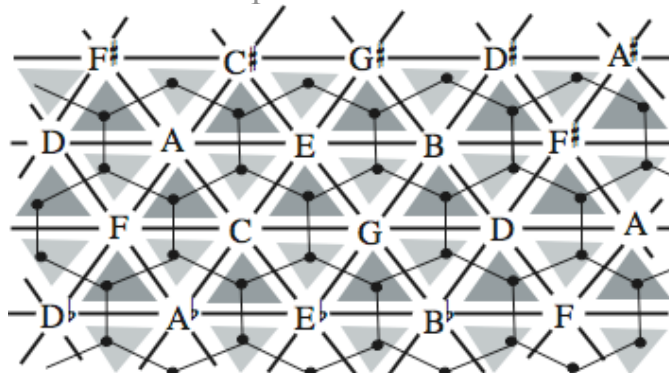



Figure 2. Triads form the dual graph of the Harmonic Network.


Pitches that belong to a given key also form compact sets of connected components with unique shapes that identify their mode and tonal center (see Figures 3(a) and 3(b)). This property was exploited in Longuet-Higgins and Steedman's shape matching algorithm for key-finding (Longuet-Higgins & Steedman 1971, Longuet-Higgins 1976). Like edit-distance for comparing strings, transformations on key shapes can be used as a metric for comparing keys, but would be less suitable for comparing objects from different hierarchical levels, for example, keys and triads.

Ftonnetz-maj.eps



(a) C major key shape.

Ftonnetz-min.eps



(b) C (harmonic) minor key shape.

Figure 3. Uniquely shaped connected components representing major and minor keys.

Most literature on the Harmonic Network alludes to the spiral structure (or toroid structure when assuming enharmonic equivalence) inherent in the grid (see Figure 4). However, the three-dimensional realization of the model is hardly used in solutions to problems of music analysis or cognition, and is not necessary for analyzing transformations on (see

Lewin 1987), or deriving group-theoretic properties of (see Balzano 1980), the network. Thus, the three-dimensional spiral configuration of the Harmonic Network is rarely used for more than illustrative purposes, and is frequently discussed only as an interesting theoretical property.




Figure 4. Spiral representation of the Harmonic Network.

Other grid models that map to cylindrical spiral structures include Lerdahl's Tonal Pitch Space (2001), based on a distance metric defined on a network of pitch classes on a cone, similar to the one discovered by Krumhansl through experimental means (1978). The resulting lattice of chord and key relationships also wrap nicely onto cylindrical spirals that fold over into tori. Another example is Shepard's double helical model for pitch relations (Shepard 1982a). These models utilize only the discrete space. Strict adherence to the lattice structure often leads only to integral values for inter-object distance produced by counting edges or transformations on the lattice.

3 Getting Inside the Spiral: Geometry of the Spiral Array

The convex hull of a set of points is the smallest shape that contains all the points. In three-dimensions, one can imagine pulling an elastic net over a set of spatial points to get their convex hull.

The Spiral Array is a geometric model that spatially represents pitches, chords and keys as points on the spiral configuration, as well as inside the spiral, of the Harmonic Network. The fundamental insight behind the model is that any collection of pitches can generate a *center of effect* (*c.e.*), that is an interior point in the convex hull of its component pitch representations, whose distance from any other element can then be measured. By using the interior space, the Spiral Array is able to represent pitches, intervals, chords (major and minor triads) and keys (major and minor) in the same spatial framework. It also represents the inter-relations among these objects as distance measured through the interior of the spiral.

The Spiral Array model begins with the spiral configuration of pitch classes as shown in Figure 4. The equation for the pitch classes is as follows:


$$\mathbf{P}(k) \stackrel{\text{def}}{=} \begin{bmatrix} x_k \\ y_k \\ z_k \end{bmatrix} = \begin{bmatrix} r \sin(k\pi/2) \\ r \cos(k\pi/2) \\ kh \end{bmatrix}, \quad (1)$$

where k marks the pitch's distance from C on the line of fifths, and C is arbitrarily set at position $[0,1,0]$.

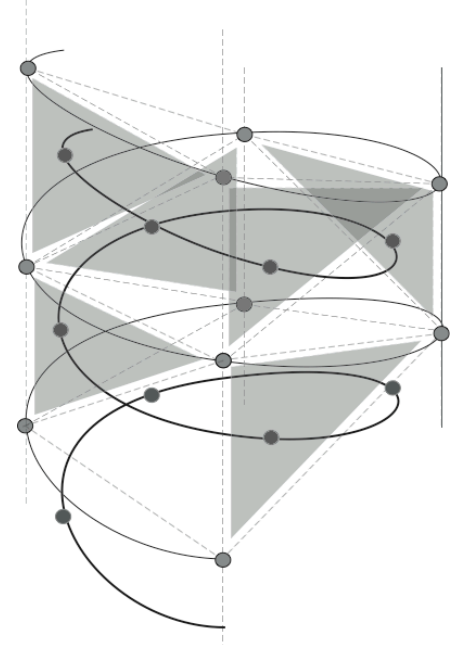
Each triad is represented as a point on the face of the triangle outlined by its component pitches. Each triad is a convex combination of its root, fifth, and third. Note that the triad representation generated in this fashion is a point in the interior of the spiral. The set of major triads form a spiral inside the pitch spiral, shown in Figure 5(a), as do the set of minor triads, shown in Figure 5(b).

The major triad equation is:

$$\mathbf{C}_M(k) \stackrel{\text{def}}{=} w_1 \cdot \mathbf{P}(k) + w_2 \cdot \mathbf{P}(k+1) + w_3 \cdot \mathbf{P}(k+4), \quad (2)$$

FECtriads1.eps

(a) Major triad spiral.

FECtriads2.eps

(b) Minor triad spiral.

Figure 5. Triad Representations in the Spiral Array.

where $w_1 \geq w_2 \geq w_3 > 0$, and $\sum_{i=1}^3 w_i = 1$. The minor triad is generated by a similar equation:

$$\mathbf{C}_m(\mathbf{k}) \stackrel{\text{def}}{=} u_1 \cdot \mathbf{P}(\mathbf{k}) + u_2 \cdot \mathbf{P}(\mathbf{k}+1) + u_3 \cdot \mathbf{P}(\mathbf{k}-3), \quad (3)$$

where $u_1 \geq u_2 \geq u_3 > 0$, and $\sum_{i=1}^3 u_i = 1$. The weights, w_i and u_i determine where on the triangle the point representing the triad resides. By choosing these weights carefully, the distance of the triad representation to its component pitch classes can reflect the desired relations among these objects. By design, the range of possible distance relations is constrained by the structure of the original Harmonic Network, as well as the way in which the triad representations are defined. For example, the constraint that the weights, w_i or u_i , should sum to one limits the point representing the triad to be inside the triangle defined by the three component pitch classes.

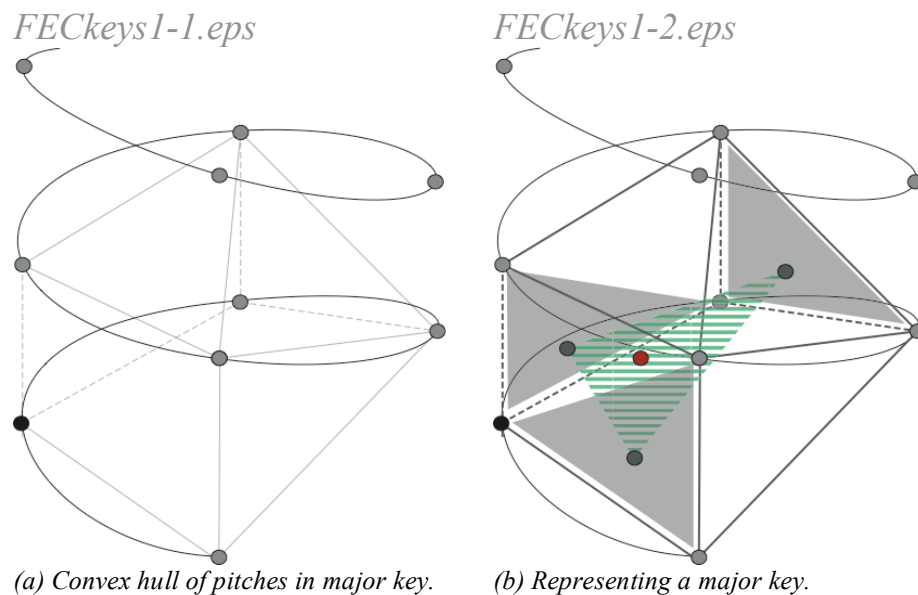


Figure 6. Major key span and representation in the Spiral Array.

As in the Harmonic Network, pitch classes belonging to a given key form compact clusters in the Spiral Array model. Figures 6 and 7 show the convex hull of the pitch classes, and the way in which the key representations are generated for the major key and the minor key respectively. The major key is represented by a spatial point in the interior of the three-dimensional spiral structures for the pitch classes, and for the major triads.

Figure 6(a) shows the convex hull of the spatial representations of the pitch classes in a given major key. Since each major key is uniquely defined by its I, V and IV triads, the major key representation is defined to be a point on the face of the triangle outlined by the spatial representations of its tonic (I), dominant (V), and subdominant (IV) triads, as shown in Figure 6(b). In Figure 6(b), the tonal center (the tonic pitch) is indicated by the black sphere, the three grey triangles and dark grey spheres in the center of the triangles represent the IV, I, and V triads. The three dark grey spheres outline an interior triangle the center of which contains the sphere that represents the major key. Figures 7(a) and (b) show the corresponding geometric objects for the (harmonic) minor key, which uses the I, V and iv triads.

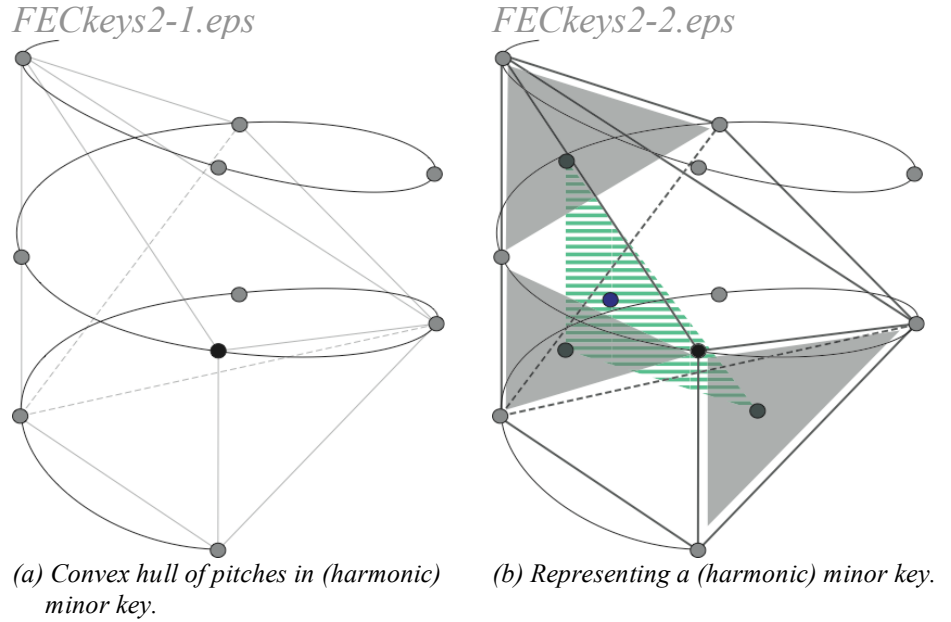


Figure 7. Minor key span and representation in the Spiral Array.

As mentioned above, the key representation is chosen to be a point on the triangle with its defining triads as vertices. Thus, the equation for the major key representation is as follows:

$$\mathbf{T}_M(\mathbf{k}) \stackrel{\text{def}}{=} \omega_1 \cdot \mathbf{C}_M(\mathbf{k}) + \omega_2 \cdot \mathbf{C}_M(\mathbf{k}+1) + \omega_3 \cdot \mathbf{C}_M(\mathbf{k}-1), \quad (4)$$

where $\omega_1 \geq \omega_2 \geq \omega_3 > 0$, and $\sum_{i=1}^3 \omega_i = 1$. The minor key definition is as follows:

$$\mathbf{T}_m(\mathbf{k}) \stackrel{\text{def}}{=} v_1 \cdot \mathbf{C}_M(\mathbf{k}) + v_2 \left[\alpha \cdot \mathbf{C}_M(\mathbf{k}+1) + (1-\alpha) \cdot \mathbf{C}_m(\mathbf{k}+1) \right] + v_3 \left[\beta \cdot \mathbf{C}_m(\mathbf{k}-1) + (1-\beta) \cdot \mathbf{C}_M(\mathbf{k}-1) \right] \quad (5)$$

where $v_1 \geq v_2 \geq v_3 > 0$, and $\sum_{i=1}^3 v_i = 1$, and $0 \leq \alpha \leq 1$, $0 \leq \beta \leq 1$.

The constraints on the weights are chosen to reflect each chord's significance in the key. For example, in Equation 4, the weight for the I chord must be no less than that for the V chord, which in turn should be

no less than that for the IV chord. In Equation 5, α and β represent the relative importance of the V versus the v triad, and of the iv versus the IV triad in the minor key. For example, when $\alpha = \beta = 1$, $T_m(k)$ represents the harmonic minor key.

Like the major and minor triads, the major and minor keys also form spiral structures, as shown in Figure 8. Figure 8(a) shows a helical sequence of major key triangles, and Figure 8(b) shows the same for the minor key triangles.

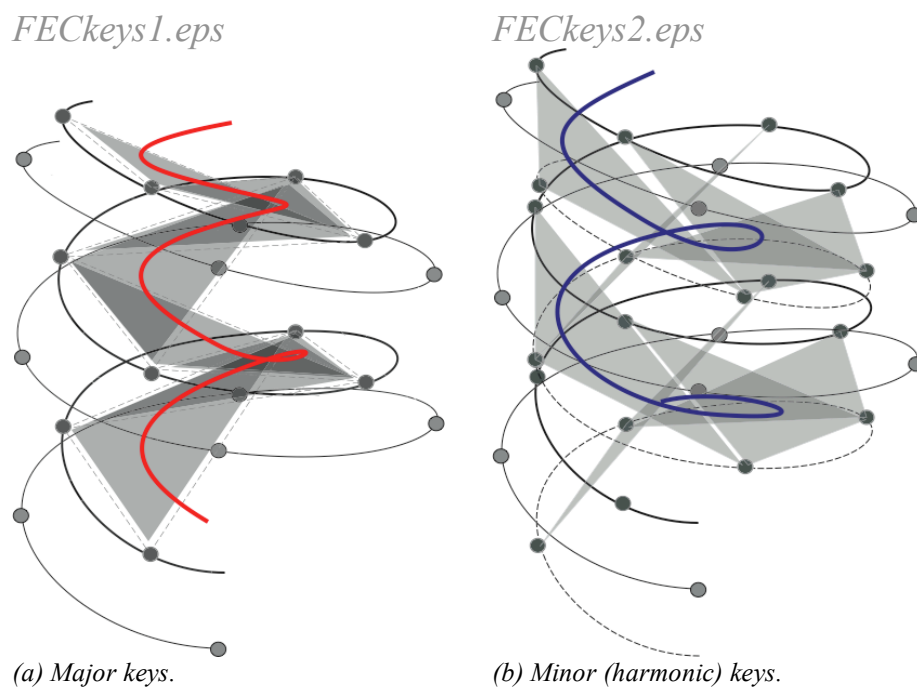


Figure 8. Key Representations in the Spiral Array.

4 A Metric to Compare Distance Between Objects

The Spiral Array model can be visualized as a set of nested spirals as shown in Figure 9. Because objects from different hierarchical levels are represented in the same space, the model can be calibrated so that the perceived distance between any two objects can be quantified.

FECribbons.eps

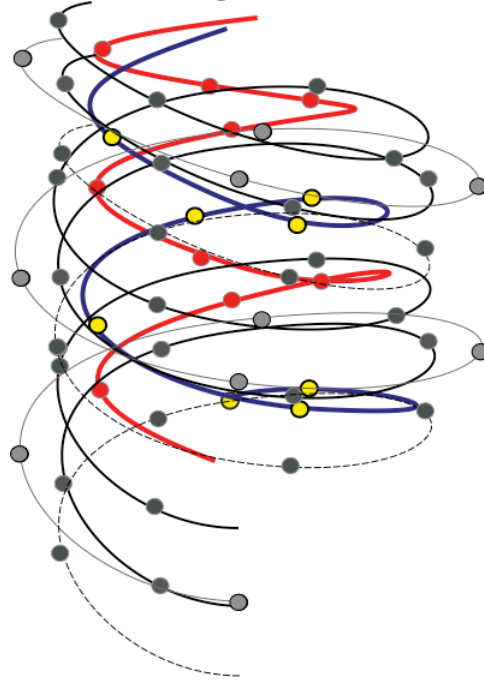
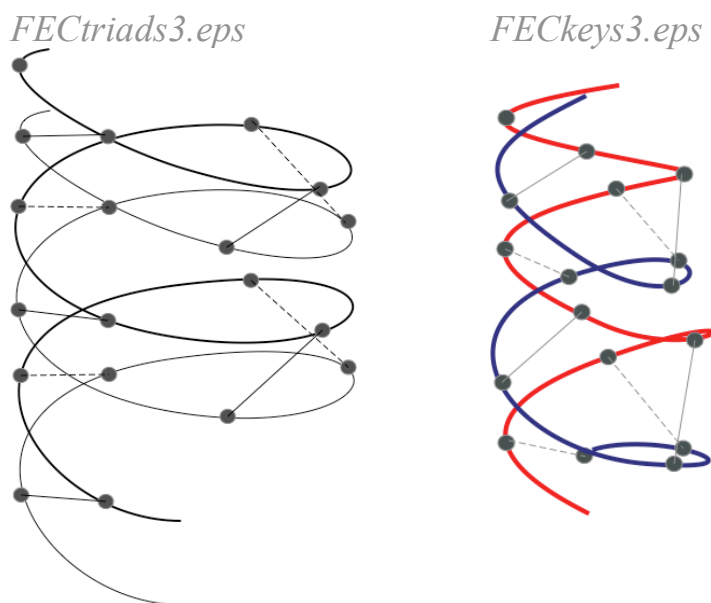


Figure 9. The Spiral Array model visualized as an array of spirals.

One of the advantages of using the interior of the spiral is that Euclidean space can now provide a metric for quantifying the distance between any two objects from any hierarchical level. For example, Figures 10(a) and (b) show the relative positions of major and minor spirals for triad representations and for key representations respectively. In these figures, traverse lines across the two helices in each pair connect objects with the same name, such as A major triad and A minor triad; the solid lines connect objects in the foreground, and the broken lines connect objects in the background.

The existence of a general metric across all hierarchical levels and for as-yet-undefined tonal objects allows one to design computer programs to compute distances between objects, and to make selections based on these distances. For example, in the Center of Effect Generator (CEG) key finding algorithm described in (Chew 2000 and 2001), pitches, say in a melody, are first mapped to their corresponding pitch class position in the pitch class spiral, and a *c.e.* of the pitch collection generated. The *c.e.* is generated by weighting each pitch by its relative duration (or some other preferred weighting scheme). The key of the melodic fragment is

then determined by finding the closest key representation on the major and minor key helices. The method generalizes to polyphonic music.



(a) Major and minor triads superimposed. (b) Major and minor keys superimposed.

Figure 10. Major and minor representations in the Spiral Array.

The CEG algorithm takes advantage of the fact that pitches, their *c.e.*, and keys are represented in the same space to recognize key using a nearest neighbor search. Figure 11 gives an illustration of a representative *c.e.* progression for a hypothetical melody, shown as a squiggly line that begins at a single pitch on the pitch class helix, and quickly winds its way closer to its key representation, indicated by a box, on the minor key spiral.

As in the interior-point methods for linear optimization, the *c.e.* does not represent the key, just as the interior points cannot be the optimal solution. By relaxing the constraint of seeking solutions only on the solution grid (the Harmonic Network or the lattice of major and minor key representations), and by moving inside the spiral(s), we allow even incomplete pitch class information, as in the case of melodies without a leading tone, to generate a *c.e.* path that leads to a key solution.

FCEG-illus.eps




Figure 11. A representative *c.e.* progression for a hypothetical melody.

The CEG method has been shown to be confounded less often, and to reach the key solution faster when compared to Longuet-Higgins & Steedman's shape matching algorithm using the Harmonic Network, and to Krumhansl's probe tone profile method, for the fugue subjects from the *Well-Tempered Clavier Book I* (Chew 2000 and 2001).

The fact that distances can be readily computed in this general space makes the algorithm particularly efficient and well suited to real-time applications. Because the computations are fast, $O(n)$, i.e. linear in the input sequence, the algorithm can be used in real-time systems (for example, Chew & François 2005 and François & Chew 2006) that identify key (and chords in a similar fashion) as the music is being played.

Similarly, the advantages of a straightforward metric, representation of general tonal objects in the same three-dimensional space, and efficiency of computation enables the design of fast algorithms for pitch spelling (see, for example, Chew & Chen 2005), and for tonal comparison and segmentation, such as in the Argus algorithm described in (Chew 2006).

5 Comparing the Resulting Geometry of the Interior-Point Approach to Krumhansl's and Lerdahl's Tonal Spaces

The previous two sections described the interior-point approach to representing higher-level tonal objects inside the pitch class spiral. This method of representing higher level objects in the same space results in an array of spirals, each representing a different type of tonal object. Since the Spiral Array inherits the spatial properties of the Harmonic Network; the incorporating of all tonal objects into the same space provides the added benefit of modeling perceived closeness among objects from different hierarchical levels.

In this section, the geometry of the Spiral Array at each hierarchical level will be compared to that of two other models, namely, Krumhansl's (1978) and Krumhansl & Kessler's (1982) multidimensional scaling solutions to pitch class and key proximity, further described in (Krumhansl 1990), and Lerdahl's Tonal Pitch Space (2001), including the pitch class, chordal, and regional space. We choose not to draw comparisons to Krumhansl et al.'s harmonic charts (described in Krumhansl 1990, p.188-212), also obtained by multidimensional scaling at this point. While these solutions consistently show that chords that function in the same key are perceived to be closer than those that do not, the exact solutions differ slightly depending on the specific key context. In this article, we consider only the representations that compare chords one to another, independent of key context. The section concludes with a comparison of the distance metrics that generate these spaces.

The purpose of these comparisons is to show that there exist direct mappings amongst the three spaces, and even when the mappings are not entirely obvious, the same tonal relations persist in the three models. While the fact that the three models exhibit striking similarities is not particularly surprising since, after all, they are modeling the same tonal phenomenon, the fact that such diverse approaches converge on similar configurations of tonal objects at each hierarchical level (pitches, triads, keys) should not be taken for granted.

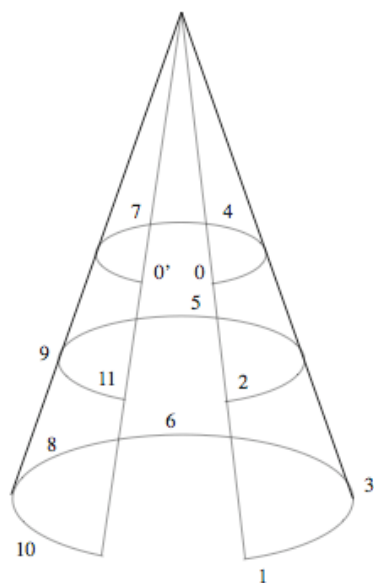
5.1 Comparison of pitch class representations

We begin with a comparison of the pitch class representations. Krumhansl's model is derived from the application of multidimensional scaling techniques to experimental data. To facilitate later comparisons

with Lerdahl's Tonal Pitch Space, an inverted version of Krumhansl's pitch cone using pitch class notation is shown in Figure 12(a). The top-most layer contains pitches in the tonic triad. The second layer contains the remainder pitches in the diatonic scale, and the base layer contains the five pitches outside of the diatonic scale.


Lerdahl's Tonal Pitch Space is based on Deutsch and Feroe's idea of hierarchically organized pitch classes (see Lerdahl 2001, p.47). Lerdahl's pitch class cone is diagrammed in Figure 12(b). The main difference between the two pitch cones is that Lerdahl's contains an additional layer highlighting the Perfect Fifth interval relation. To map the layers from Krumhansl's to Lerdahl's model, one simply adds the missing layer and allows pitch classes at each level in Krumhansl's pitch cone to carry over to the next layer down the cone.

Fcone-ck.eps



(a) Krumhansl's pitch cone (inverted).

Fcone-fl.eps



(b) Lerdahl's pitch class cone.

Figure 12. Pitch Class Cones.

As can be observed in Figure 12, the pitch arrangements in the Spiral Array mirror that of Krumhansl's and Lerdahl's pitch class cones. Figure 13 shows the Spiral Array pitch classes that correspond to each layer in Krumhansl's cone – the tonic (corresponding to pitch class 0 in the cones) is shown as a white sphere in the figures other than the first,

Figure 13(a), as a reference although it is not one of the pitches in the layers shown in Figures 13(b) and (c). In Figure 13(d), some of the pitch classes higher up in the helix are not visible in the diagram. Figure 14 shows the Spiral Array pitch classes that correspond to each layer in Lerdahl's pitch class cone.

Observe that the hierarchical ordering of the distance from the tonic to each level of pitches in the cone representation is part of the Spiral Array structure. Compare Figures 13 and 12(a). Each layer that is closer to the base of the cone contains pitches that map to positions on the Spiral Array that are progressively farther away from the tonic. If one traced a line through the small spheres, in the sequence indicated by the corresponding level in the pitch class cone, in Figures 13(b) and 13(c), one would draw Dali clock-like shapes that wrap around ever widening spheres of influence away from the tonal center.

Compare Figures 14 and Figure 12(b). Each layer closer to the base of the cone maps to positions on the Spiral Array that defines an expanding compact set. In this case, the convex hulls of the pitch class sets grow in size from Figures 14(a) through 14(d).




Figure 13. From the Spiral Array to Krumhansl's Pitch Class Cone.

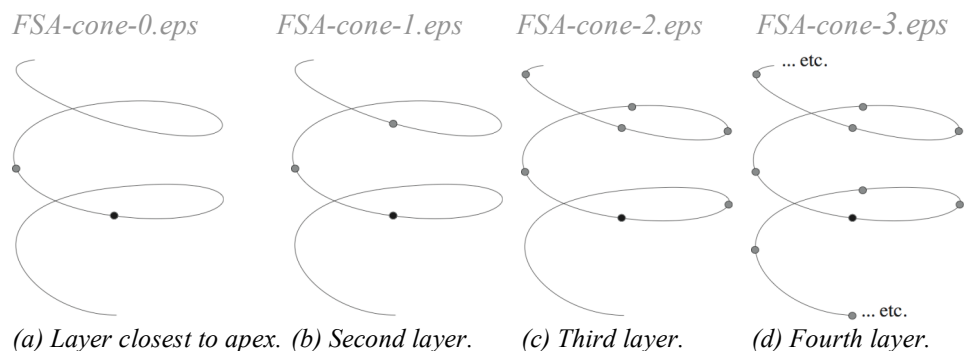


Figure 14. From the Spiral Array to Lerdahl's Pitch Class Cone.

5.3 Comparison of chord representations

In this section, we show the correspondence between Lerdahl's chordal space and the triad structures in the Spiral Array model. Lerdahl defines the local chord distance as the sum of the number of shifts along the circle-of-fifths necessary to transform one chord to the other, and the number of distinct pitch classes between the chords being compared (Lerdahl 2001, p.55–57). The table in the lower part of Figure 15 shows Lerdahl's chordal space, a spatial arrangement of chords in a key that mirrors the proximity relations derived using Lerdahl's distance metric.

vii ^o	ii	IV	vi	I	iii	V
iii	V	vii ^o	ii	IV	vi	I
vi	I	iii	V	vii ^o	ii	IV
ii	IV	vi	I	iii	V	vii ^o
V	vii ^o	ii	IV	vi	I	iii
I	iii	V	vii ^o	ii	IV	vi
IV	vi	I	iii	V	vii ^o	ii

Figure 15. Relations among chord functions in Lerdahl's Chordal Space.

Ftonnetz1-cim.eps

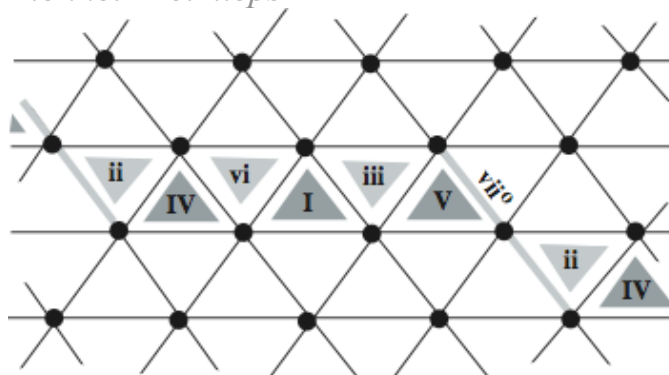


Figure 16. Lerdahl's Chord Space mapped to the Harmonic Network.

Each row in the chordal space, shown in Figure 15, traces the path highlighted in the Harmonic Network, as shown in Figure 16. The sequence V, I, IV is indicated by the dark gray triangles, followed by the gray line representing the vii° chord, and the iii, vi, ii sequence indicated by the light gray triangles. As in Lerdahl's chordal space, iii and vi neighbor I and V, and IV and I respectively, and vii° is next to V and ii.

Because the Spiral Array inherits the pitch relations shown in the Harmonic Network, the same chord relations indicated above can be demonstrated in the Spiral Array. The highlighted chords in the Harmonic Network in Figure 16 can be thought of as a sequence of alternating major and minor triads, connected by the vii° chord. This same pattern is evident in the Spiral Array, as shown in Figure 17. In Figure 17, the broken lines mark the edges of the major and minor triad triangles, and the angled gray bars indicate the location of the vii° chord. The black wavy line connects the alternating major and minor chord sequence and the vii° chord, showing the parallel with the Harmonic Network, and thus the connection with Lerdahl's chordal space.


FFLchords spiral.eps

Figure 17. Lerdahl's Chordal Space mapped to the Spiral Array.

5.4 Comparison of key representations

This section compares Lerdahl's regional (key) space, Krumhansl & Kessler's multidimensional scaling solution for the 24 major and minor keys, and the key representations in the Spiral Array. Lerdahl's regional space is built the same way as his chordal space, by finding the nearest local tonic chords, according to the general chord distance metric described in the previous section (Lerdahl 2001, p.59–65). The general chord distance metric takes into account the key context of a chord, and is the sum of the distance, on the circle of fifths, between the two key contexts, that of the two chords, and the number of distinct pitch classes. The resulting table of key relations is shown on the right half of Figure 18, not according to scale as all neighbors should be equidistant.

FECkeys4.eps

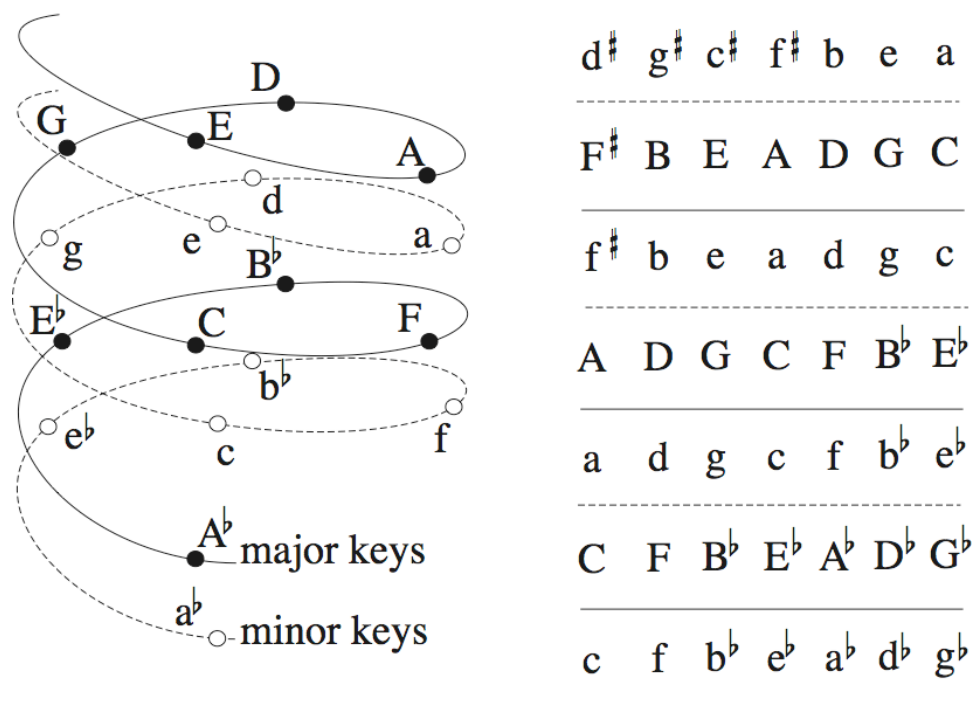


Figure 18. Key Representations in the Spiral Array and Lerdahl's Regional Space.

Lerdahl's regional space parallels Krumhansl and Kessler's (1982) chart of the multidimensional scaling solution for key relations. In (Krumhansl 1990) and the original 1982 article, the chart shows the same configuration of major and minor keys as that in the regional space, with a very slight shift in alignment between the major key and minor key diagonals.

The connection between Lerdahl's regional space, Krumhansl and Kessler's key chart, and the Spiral Array is made more apparent by the consideration of the chromatic pitch set. As can be seen in Figure 18, the key relations represented in the different models are equivalent. The two helices on the left part of Figure 18 show the major key and minor key spirals, shown as solid line and broken line helices respectively. For direct comparison, the solid and broken lines in the regional space / key chart on the right highlight the parallels between the two models.

Assuming enharmonic equivalence, the spiral structures would wrap around to form a torus, just like Lerdahl's regional space and Krumhansl

and Kessler's key chart. Findings on mental models of key relations have been remarkably consistent. For example, more recent experiments using Magnetic Resonance Imaging of brain activity with subjects listening to melodies in different keys have further confirmed the toroid structure of key relations (Zatorre & Krumhansl 2002).

5.4 Remarks on the distance metrics

While we have focused primarily on the similarities amongst the models obtained by the three approaches thus far, it is useful to now highlight the specific differences in the metric distances employed by these approaches.

The Krumhansl pitch class cone and the Krumhansl & Kessler key charts are obtained from multidimensional scaling of listener ratings, and is generated entirely from experimentally obtained data. In contrast, both Lerdahl's Tonal Pitch Space and the Spiral Array first defines a model, and methods for computing distances amongst tonal objects, before setting out to ascertain the spatial organization of these objects.

Lerdahl's method for computing pitch class distance is a combination of horizontal and vertical distance in the basic pitch class space. The method for computing chord and key distances is a combination of counting discrete steps on the circle-of-fifths and the number of distinct pitch classes. Specific directions for computing distances amongst chords, amongst chords with respect to their tonal contexts, and amongst key contexts differ one from another.

As in multidimensional scaling, Lerdahl's spatial organization of the objects depends on the distances between objects. The difference from Krumhansl and Krumhansl & Kessler's approach is that one obtains distance ratings from listeners, while the other computes distances according to specific rules. The method by which these distances are translated into spatial arrangements also differs.

In the Spiral Array, objects are computed as weighted averages of their components, and represented as spatial points in the structure. A distinct advantage of the interior-point approach is that objects from different hierarchical levels are represented in the same space, thus facilitating inter-level object comparisons. Using this approach, one can define distances between, say, a pitch class and a key, or a key and a particular chord. The approach treats all objects equally in the same three-dimensional space, thus blurring the boundaries between pitches, chords,

and keys, and other objects one might create in this space, such as pitch sets, sets of pitch sets, and so on.

The interior-point approach has the added benefit of low computational complexity, which becomes an asset when analyzing large numbers of notes, or when considering more complex pitch class sets. All objects in the Spiral Array are generated by one simple rule, that of computing the *c.e.* of the component parts. In the Spiral Array, the constraints of a discrete space are relaxed to create a continuous interior space that can be used to solve problems of tonal recognition. Combinatorial problems in discrete space are transformed into computationally simpler ones, such as nearest-neighbor searches, in continuous space.

The continuous space in the Spiral Array also allows incomplete or noisy pitch sets to define objects within it, and for tonal recognition problems to work well in spite of such challenges. In a sense, the interior-point approach is situated between strictly discrete treatment of tonal space, and a probabilistic treatment of the same space. The entire approach is defined mathematically, without employing experimental data.

Note that the interior-point approach is not limited to the Harmonic Network. The idea extends to any geometric model that clusters objects that form higher level structures.

6 Conclusions

This paper posits and validates the use of the interior-point approach to modeling higher-level structures using the spiral configuration of the Harmonic Network. To demonstrate this concept, a geometric interpretation of the Spiral Array model and its fundamental idea of representing higher-level objects in the interior as convex combinations of the representations of the lower level components has been presented. This interior-point approach not only preserves the pitch relations of the original lattice, it also generates higher-level structures that are consistent with other researchers' results. In particular, mappings were shown among Lerdahl's Tonal Pitch Space, Krumhansl's and Krumhansl & Kessler's spatial representations of pitch class and key relations, and the Spiral Array model. Arguments for the advantage of employing the interior-point approach have been provided, including its simplicity, computational efficiency, robustness, and generality.

Acknowledgments

I thank Jeanne Bamberger for our numerous and stimulating discussions on tonality, and for introducing me to the work of Christopher Longuet-Higgins and Mark Steedman, and of David Lewin. This material is based upon my dissertation work supported by a Josephine de Kármán Dissertation Fellowship, and ensuing research supported by the National Science Foundation (NSF) Grant No. 0347988, and a fellowship at the Radcliffe Institute for Advanced Study. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the Josephine de Kármán Foundation, NSF, or the Radcliffe Institute for Advanced Study.

References

- Balzano, Gerald J. (1980). "The group-theoretic description of 12-fold and microtonal pitch systems," *Computer Music Journal*. (4:4), 66–84.
- Chew, Elaine (2000). *Towards a Mathematical Model of Tonality*, Ph.D. dissertation, Operations Research Center, MIT, Cambridge, MA.
- Chew, Elaine (2001). "Modeling Tonality: Applications to Music Cognition," in *Proceedings of the 23rd Annual Meeting of the Cognitive Science Society*, ed. Johanna D. Moore and Keith Stenning (Mahwah, NJ: Lawrence Erlbaum Associates), 206–211.
- Chew, Elaine (2006). "Regards on Two Regards by Messiaen: post-tonal music segmentation using pitch context distances in the Spiral Array," *Journal of New Music Research*, Special Issue on Sound and Music Computing, selected papers from the First Sound and Music Computing Conference, revised and expanded, ed. Gérard Assayag. (34:4), 341–354.
- Chew, Elaine and Chen, Yun-Cheng (2005). "Real-Time Pitch Spelling Using the Spiral Array", *Computer Music Journal*. (29:2), 61–76.
- Chew, Elaine and François, Alexandre R. J. (2005). "Interactive multi-scale visualizations of tonal evolution in MuSA.RT Opus 2", *Computers in Entertainment*, (3:4), article 3E, 16 pages.
- Cohn, Richard (1996). "Maximally smooth cycles, hexatonic systems, and the analysis of late-romantic triadic progressions," *Music Analysis*. (15), 9–40.
- Cohn, Richard (1997). "Neo-riemannian operations, parsimonious trichords, and their tonnetz representations," *Journal of Music Theory*. (41), 1–66.
- Cohn, R. (1998). "Introduction to Neo-Riemannian Theory: A Survey and a Historical Perspective," *Journal of Music Theory*. (42), 167–180.
- Dantzig, George B. (1963). *Linear Programming and Extensions* (Princeton, NJ: Princeton University Press).
- François, Alexandre R. J. and Chew, Elaine (2006). "An Architectural Framework for Interactive Music Systems", in *Proceedings of the*

- International Conference on New Interfaces for Musical Expression*, 150-155.
- Hillier, F. S., and Lieberman, G. J. (2001). *Introduction to Operations Research*, Seventh Edition (New York, NY: McGraw-Hill Higher Education).
- Karmarkar, N. (1984). "A new polynomial-time algorithm for linear programming," *Combinatorica*. (4), 373–395.
- Krumhansl, Carol L. (1978). *The Psychological Representation of Musical Pitch in a Tonal Context*, Ph.D. dissertation, Stanford University, Stanford, CA.
- Krumhansl, C. L. and E. J. Kessler (1982). "Tracing the dynamic changes in perceived tonal organisation in a spatial representation of musical keys," *Psychological Review*. (89), 334–368.
- Krumhansl, Carol L. (1990). *Cognitive Foundations of Musical Pitch* (New York, NY: Oxford University Press).
- Lerdahl, Fred (2001). *Tonal Pitch Space* (Cambridge, MA: MIT Press).
- Lewin, David (1982). "A formal theory of generalized tonal functions," *Journal of Music Theory*. (26), 32–60.
- Lewin, David (1987). *Generalized Musical Intervals and Transformations* (New Haven, CT: Yale University Press).
- Longuet-Higgins, H. Christopher (1976). "Perception of melodies," *Nature*. (263), 646–653.
- Longuet-Higgins, H. Christopher and Mark J. Steedman (1971). "On interpreting Bach," in *Machine Intelligence*, ed. Bernard Meltzer and Donald Michie (Edinburgh, UK: Edinburgh University Press). (6), 221–242.
- Shepard, Roger N. (1982). "Structural representations of musical pitch," in *The Psychology of Music*, ed. D. Deutsch (New York, NY: Academic Press), 335–353.
- Zatorre, Roger J. and Carol L. Krumhansl (2002). "Mental models and musical minds," *Science*. (298), 2138–2139.