Topology of Musical Data

William A. Sethares
Department of Electrical and Computer Engineering,
University of Wisconsin, Madison, USA, sethares@ece.wisc.edu

November 27, 2010

Abstract

Techniques for discovering topological structures in large data sets are now becoming practical. This paper argues why the musical realm is a particularly promising arena in which to expect to find nontrivial topological features. The analysis is able to recover three important topological features in music: the circle of notes, the circle of fifths, and the rhythmic repetition of timelines, often pictured in the necklace notation. Applications to folk music (in the form of standard MIDI files) are presented and the bar codes show a variety of interesting features, some of which can be easily interpreted.

1 Introduction

Carlsson and his coworkers [3] have recently introduced a way of parsing large data sets using *barcodes* that show the Betti numbers of an underlying topological space. A parameter ϵ is slowly increased and a cloud of points (which may be pictured in $\mathbb{R}^n$) are identified whenever they are closer together than ϵ. For ϵ small, the structure is trivial: all points are separated from all other points. For large ϵ, the structure is again trivial, all points are identified together and the space is topologically equivalent to a single point. In between small and large, cycles may appear in two dimensions, spheroids in three dimensions, and other higher dimensional analogs may appear and disappear as ϵ changes. Features which persist over a range of ϵ are called persistent, and likely reflect some underlying structure in the data. This technique is called persistent homology, and a number of applications have begun to appear in areas such as image processing [6] and in the analysis of biological data [14].

The search for such topological features in data is only beginning, and it makes sense to look in places where one may reasonably expect to find interesting structures. One such place is in musical data. It is a commonality that in musical scales there is a "circle of notes" which may be pictured as in Figure 1. This figure encodes two important aspects of musical perception: notes that are near each other in frequency (such as C and $C\sharp$) are perceptually close, and notes that are an octave apart (such as low C and high C) are perceptually close, even though they differ by a factor of two in frequency. Obviously, Figure 1 shows a circle, and it is reasonable to ask if such circles can be recovered from an analysis of musical data.

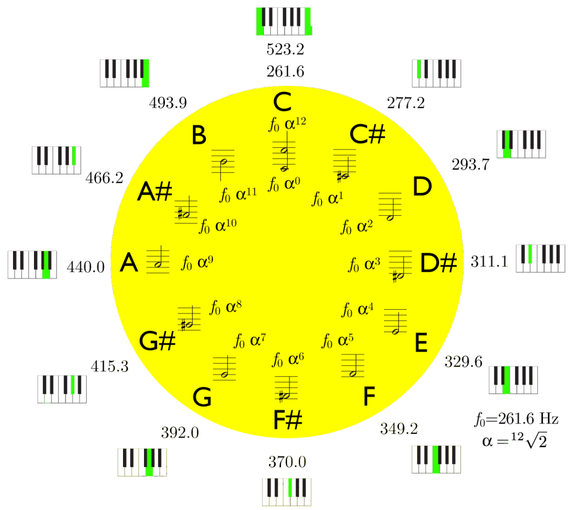

Figure 1: A surprising number of insights about musical structure are displayed in this "circle of notes," which is like a clock face on which the hours of the day have been replaced by note names.

A second well known topological structure in music theory is the "circle of fifths," shown here in Figure 2, which is taken from the Wikipedia article of the same name [16]. The circle of fifths is a standard way musicians and composers talk about the close relationships between musical scales and keys, and represents another way of interpreting the distance between musical chords and scales. Again, Figure 2 shows a circle, and it is reasonable to ask if this circle can be recovered from an analysis of musical data.

Representing temporal cycles as spatial circles is an old idea: Safî al-Din al-Urmawî, the 13th century theoretician from Baghdad, represents both musical and natural rhythms in a circular notation in the *Book of Cycles* [1]. Time moves around the circle (usually in a clockwise direction) and events are depicted along the periphery. Since the "end" of the circle is also the "beginning," this emphasizes the repetition inherent in rhythmic patterns. Anku [2] argues that African music is perceived in a circular (rather than linear) fashion that makes the necklace notation, shown here in Figure 3, particularly appropriate. Clearly, Figure 3 shows a circle, and it is reasonable to ask if this circle can be recovered from an analysis of musical data.

This paper shows that all three of these circles can indeed be located in the barcodes drawn using the techniques of persistent homology.

There are several different types of musical data [12], including

1. spectral data that shows the internal structure of individual sounds,

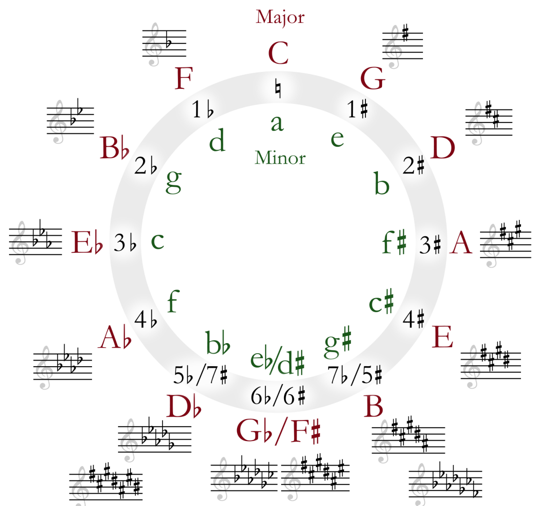

Figure 2: The circle of fifths shows the relationships among the tones of the equal tempered chromatic scale, standard key signatures, and the major and minor keys.

2. audio data that provides a literal (numerical) representation of sound waves,
3. symbolic note data (such as occurs in a musical score or in standard MIDI files) that provides instructions for a performer to play a specific piece of music,
4. analytical data (such as Roman numeral analysis [10] or Schenkerian analysis [11]) that can be derived from musical scores to provide theorists with tools to understand musical progressions, and
5. periodic data that relates to the temporal and rhythmic aspects of music [13].

Any kind of periodicity, if properly embedded, may lead to nontrivial topological structures, and periodicities may occur in rhythmic data at several levels in a metric hierarchy. In principle, it may be possible to locate circles (of notes, of fifths, of periodicities in rhythmic or tonal material) or other topological structures in many of the above kinds of data. For this initial investigation, we use MIDI note data, since this is the level at which the circles in Figures 1-3 are conceptualized, and hence the level at which success is most likely.

Section 2 defines a *pitch-class metric* that allows a clear display of the circle of notes. Data is taken from a standard repertoire of traditional folk melodies, and the barcodes are shown to display information about the musical scale used in the piece (using the Betti$_0$ barcode) and to reflect the circle of notes itself in the Betti$_1$ barcode. More generally, the note data can be gathered into higher dimensional time-delay embeddings. When this is done for the same folk melodies, interesting new (and as yet unexplained) features arise in the barcodes.

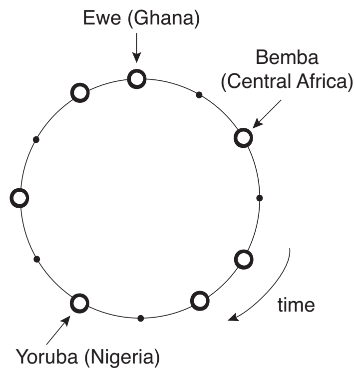

Figure 3: Traditional rhythms of the Ewe (from Ghana), the Yoruba (from Nigeria) and the Bemba (from Central Africa) are all variants of the "standard rhythm pattern" described by King [5]. These timelines are represented in the "necklace notation" as having different starting points.

Section 3 looks at the issue of measuring chord and scale data (i.e., when multiple notes occur simultaneously) using the techniques of persistent homology. A metric called the *chord-class metric* allows a clear display of the circle of fifths on specially chosen synthetic data. Chord data is then taken from a traditional folk melody and MIDI data is used to analyze a collection of four-voice Bach chorales. The barcodes are shown to display information about the musical chords used in the piece. More generally, the chord data can be gathered into higher dimensional time-delay embeddings, and again new and unexplained patterns emerge in the barcodes.

Section 4 analyzes some simple rhythmic patterns (such as those of Figure 3) using the techniques of persistent homology. The appropriate metric is analogous to those above, though it must be modified to reflect the fact that temporal data is not perceived in a logarithmic fashion. Once again, the circular structures are immediately evident from the barcodes.

2 Melodic Barcodes and the Circle of Notes

The pitches of musical tones are generally perceived as a function of frequency in a logarithmic fashion. Thus the 15.6 Hz "distance" from C to $C\sharp$ is perceived to be the same size as the 24.7 Hz "distance" from $G\sharp$ to A (refer to Figure 1 for the origin of these numerical values). Accordingly, it is reasonable to consider a measure that operates on $\log$ frequency rather than on frequency itself. A metric like $|\log_2(f) - \log_2(g)|$ captures this along with the notion that nearby tones (with fundamental frequencies f and g) on the circle of notes should have a small numerical distance. But this measure fails to capture the idea that the C at 261.6 Hz and the high C at 523.2 Hz are effectively the same. This is what happens when a man sings along with a woman (or when a woman sings along with a child): the "same" note is actually a factor of two apart in frequency. This second notion of closeness can be incorporated by using

$$s = \text{ mod } (|\log_2(f) - \log_2(g)|, 1) \tag{1}$$

but this is not a metric since it fails the triangle inequality. It can, however, be made a metric by defining the distance between two notes f and g, expressed in terms of their fundamental frequency in Hz, as

$$d(f,g) = \min(s, 1-s) \tag{2}$$

where s is from (1). This measure may be interpreted as a measure of distance between "pitch classes" [9], since it identifies all Cs, all $C\sharp$s, etc. into equivalence classes. We call (1)-(2) the *pitch-class* metric.

To verify that this metric makes sense, consider a major scale consisting of the eight notes C, D, E, F, G, A, B, C with the frequencies as specified in Figure 1. The `Plex` software [8] is designed to "calculate the persistent homology of finite simplicial complexes... generated from point cloud data." In this case, the point cloud is defined by a matrix of distances between all pairs of the eight notes using the pitch-class metric. The resulting barcodes are shown in Figure 4.

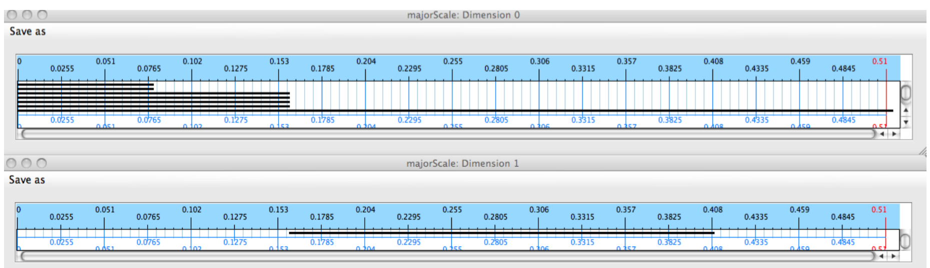

Figure 4: Barcodes calculated by the Plex software show the number of connected components in the Dimension 0 plot (top) and the number of circles in the Dimension 1 plot (bottom), as the size parameter ϵ varies.

These two plots are straightforward to interpret. When the size parameter ϵ is small, there are seven distinct notes. Though we input eight notes, the high C has exactly the same distances to all the other notes as the low C under the pitch-class metric (1)-(2), and thus the barcode merges these two tones even at $\epsilon = 0$. When ϵ reaches 0.08, the two half steps (the intervals between E-F and B-C) merge. When ϵ reaches 0.16, the five remaining connected components (all the major seconds) merge into one. Thus $\text{Betti}_0 = 1$ for all greater ϵ. At $\epsilon = 0.16$, the Dimension 1 code shows a single component, which persists until $\epsilon = 0.4$. This $\text{Betti}_1 = 1$ feature is exactly the "circle of notes" shown in Figure 1.

The example of Figure 4 was built specifically with the circle of notes in mind, so it is perhaps unsurprising that the circle appears. Will such shapes appear in real music? The website [7] contains a large selection of traditional melodies, with most tunes available in both sheet music and as standard MIDI files. The musical score for

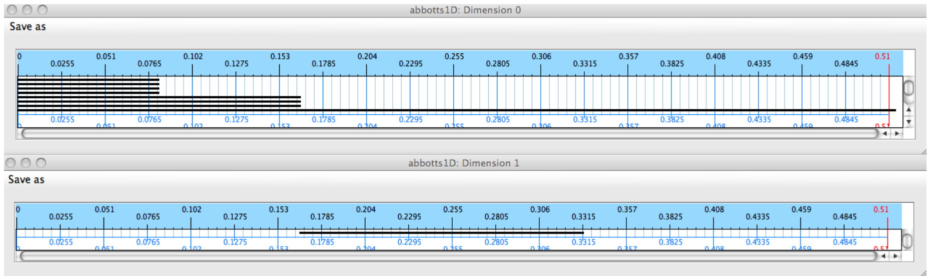

Figure 5: Barcodes for the traditional folk tune "Abbott's Bromley Horn Dance" (see Figure 6) show many of the same features as the major scale barcodes of Figure 4. The distribution of whole and half steps are clear from the Betti$_0$ code for small ϵ while the circle of notes appears again in Betti$_1$ when $0.16 < \epsilon < 0.33$.

"Abbott's Bromley Horn Dance" is shown in Figure 6 and the corresponding barcodes are shown in Figure 5.

The top barcode in Figure 5 shows eight lines, which correspond to the eight notes that appear in the score (observe again the insensitivity to octave). Four disappear at $\epsilon = 0.08$, which correspond to the four half steps ($F\sharp$-F, $D\sharp$-E, B-C, and D-$D\sharp$). Three more disappear at $\epsilon = 0.16$. Along with the constant bar, these correspond to the four whole steps (E-$F\sharp$, G-A, A-B, and C-D). All of these join into one bar for all larger ϵ. The region $0.16 < \epsilon < 0.33$ is characterized by Betti$_0 = 1$ (one connected component) and Betti$_1 = 1$, one circle. This is again the circle of notes. In fact, all the melodies from the website [7] show this same structure, though the number of half and whole steps changes to reflect the scale of the piece, and the exact extent of the Betti$_0$=Betti$_1 = 1$ region is somewhat variable.

The analyses of Figures 4 and 5 may be somewhat naive because they suppresses temporal information in the melody. This can be addressed by using a time-delay embedding, which is a common procedure in time series analysis. Suppose that a melody consists of a sequence of notes with fundamentals at $f_1, f_2, f_3, f_4 \ldots$. These may be combined into pairs (a two-dimensional time-delay embedding) by creating the sequence (f_0, f_1), (f_1, f_2), (f_2, f_3), ... The distances between such pairs can be calculated by adding the distances between the notes element-wise using the pitch-class metric.[1] Building a matrix of all such distances for "Abbott's Bromley Horn Dance" and calculating the barcodes gives Figure 7.

The Dimension 0 barcodes (the top plot in Figure 7) can be interpreted as showing the distances between pairs of notes as the melody progresses over time. Thus there are 11 pairs of notes that are at a distance of one-half step, since 11 lines end at $\epsilon = 0.08$. There are 19 pairs that differ by a whole step since 19 lines end at $\epsilon = 0.16$. Above this value, all pairs have merged into a single connected component. This can be interpreted as saying that the melody progresses primarily by stepwise motion, and that no pairs of tones are isolated from any other pairs of tones (though of course there are many

[1] Distances may alternatively be calculated using the chord-class metric d_{cc} of (3). The differences in the resulting barcodes appear to be subtle, at least for the traditional melodies of [7].

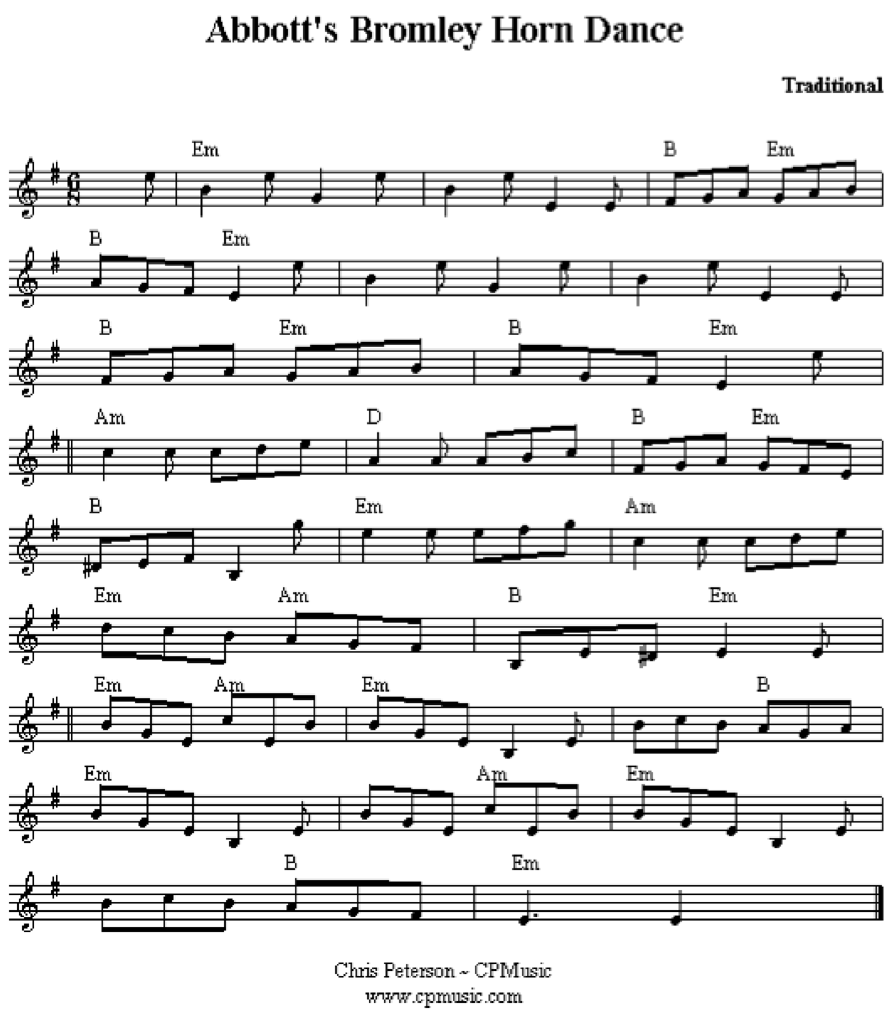

Figure 6: The traditional melody "Abbott's Bromley Horn Dance" is taken from Chris Peterson's collection [7]. The standard MIDI version of this melody is analyzed using the ideas of persistent homology in Figures 5, 7, 8, 10, and 11, and a barcode analysis of the rhythm is shown in Figure 16.

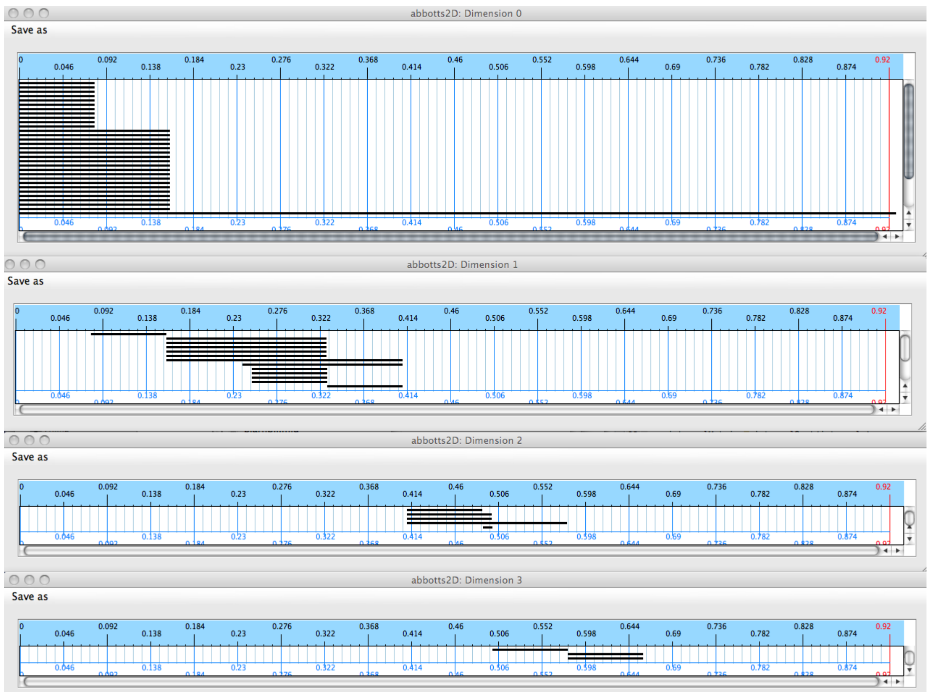

Figure 7: Barcodes for the two-dimensional time-delay embedding of “Abbott’s Bromley Horn Dance” (see Figure 6) are considerably more interesting than those in Figure 5.

individual pairs with larger distances).

The Dimension 1 barcodes (the second plot in Figure 7) shows the number of circles present at each value of ϵ, the Dimension 2 barcodes (the third plot) show the number of hollow spheres as a function of ϵ, and the Dimension 3 barcodes (the bottom plot) show the distribution of 4D holes in 4-space (do these have a name?). It would be interesting to try and interpret these higher dimensional structures in terms of the underlying piece of music. Other musical pieces (from the same library) have qualitatively similar structures, though the details appear to differ in intriguing ways.

Longer temporal information can be incorporated by using longer time-delay embeddings. If a melody consists of a sequence of notes with fundamentals at f_1, f_2, f_3, f_4..., these can be combined into triplets (a three-dimensional time-delay embedding) by creating the sequence (f_0, f_1, f_2), (f_1, f_2, f_3), $(f_2, f_3, f_4), ...$ The distances between such triplets can be calculated by adding the distances between the notes elementwise using the pitch-class metric. Building a matrix of all such distances for "Abbott's Bromley Horn Dance" and calculating the barcodes gives Figure 8.

Again, it is straightforward to interpret the Dimension 0 barcodes as distances between triplets of notes in the melody. The higher dimensional structures are again somewhat enigmatic, though presumably they indicate something about the pieces being analyzed.

3 Harmonic Barcodes and the Circle of Fifths

In order to look for the second major topological feature that should exist in musical data (the circle of fifths of Figure 2) it is necessary to generalize the metric to consider scalar harmony, multiple pitches considered simultaneously. Perhaps the most straightforward generalization of the pitch-class metric is to add the (pitch-class) distances between all elements of the vectors, as was done for the time-delay embeddings of the previous section. This metric distinguishes chord inversions: for instance, a C major chord in root position (C-E-G) would be distant from a C major chord in third position (G-C-E). While this is desirable in some musical situations, it is undesirable when looking for structures that involve musical key, where (say) all C major chords are identified irrespective of inversion and all C major scales are identified irrespective of the order in which the pitches are listed. For example, the ascending C major scale and the descending C major scale are both the same entity, and the metric should reflect this realm of musical perception. In terms of the levels of musical data (items (1)-(5) on page 2), pitch-classes are appropriate for the symbolic note level data (item 3) while the circle of fifths lies at the analytical level (item 4).

Accordingly, let $f = (f_1, f_2, ...f_n)$ and $g = (g_1, g_2, ...g_n)$ be two n-tuples, and define the distance

$$d_{cc}(f, g) = \min_P d(f, Pg) \tag{3}$$

where P ranges over all possible permutation matrices and where $d(\cdot, \cdot)$ is the pitch-class metric of (1)-(2) applied in an element-by-element fashion. This *chord-class metric* calculates the (elementwise) pitch-class distance between f and all the permutations

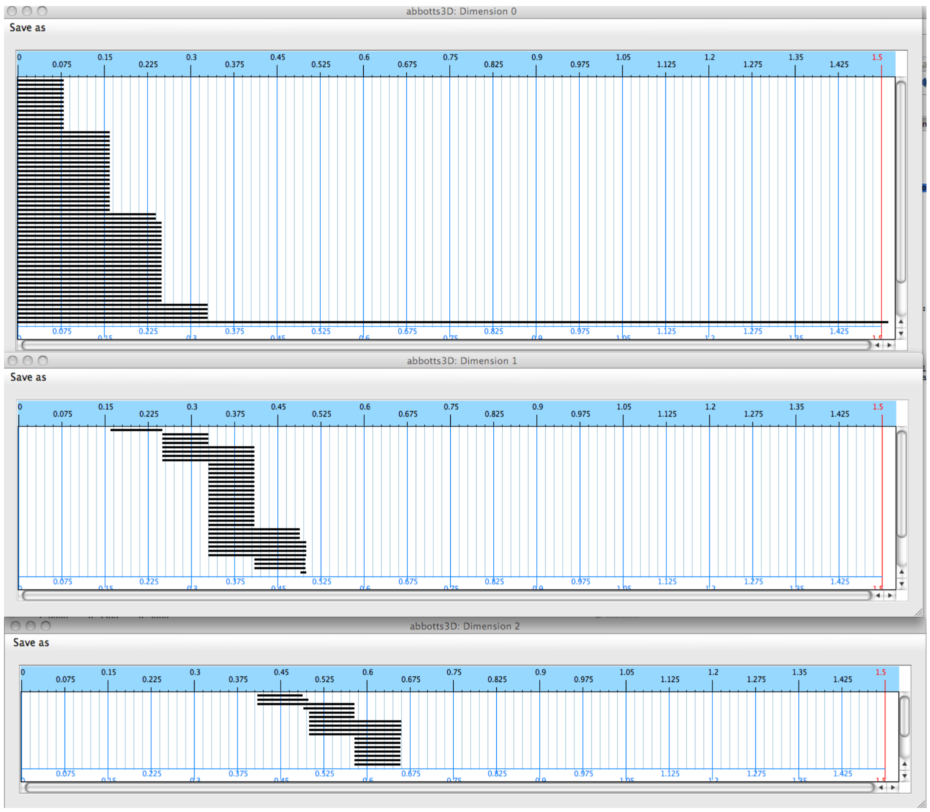

Figure 8: Barcodes for the three-dimensional time-delay embedding of "Abbott's Bromley Horn Dance" (see Figure 6) show a remarkable array of features that would be great to understand. While the dimension 0 plot is straightforward (showing the number of melodic triplets at each distance), the Dimension 1 and 2 plots contain a fascinating collection of circles and higher dimensional analogs that persist over a non-trivial range of ϵ.

of g and hence is invariant with respect to chord and scale inversion; all reorderings of the elements of f and g are placed in the same equivalence class.

To verify that this metric makes sense, consider a progression that moves around the circle of fifths: C major to G major to D major etc, all the way back to F and finally C. Inputting these seven-note sets into Plex and calculating the barcodes gives Figure 9. Under this metric, scales that are a fifth apart (such as C major and G major) have a distance of 0.08 and this explains the twelve lines that merge down to a single connected set at $\epsilon = 0.08$ in the Dimension 0 (top) plot. For $0.08 < \epsilon < 0.33$, the Dimension 1 barcode shows a single persistent bar; this is the circle of fifths! There are also some higher dimensional features for larger ϵ, but the exact meaning of these is not clear.

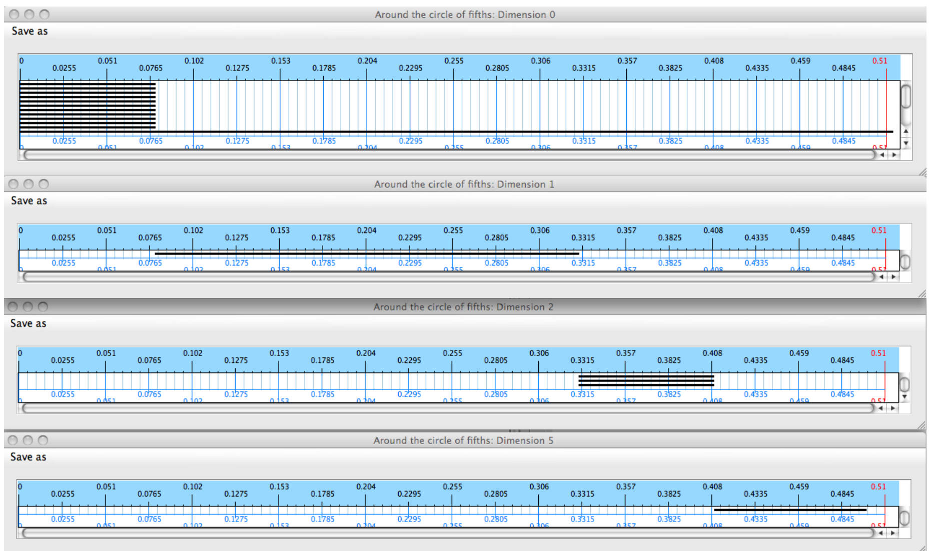

Figure 9: Barcodes for a chord progression consisting of one cycle around the circle of fifths. The circle of fifths is the persistent line from $0.08 < \epsilon < 0.33$ in the Dimension 1 plot.

The same metric can, of course, be applied to chord progressions. The chords from the score of "Abbott's Bromley Horn Dance" in Figure 6 were entered manually, and the barcodes calculated in Figure 10. Since there are only four chords (Em, B, Am, D), there is not much structure. The Dimension 0 barcode shows the distances between the four chords, and the dimension 1 barcode only shows a circle for $0.33 < \epsilon < 0.4$. It is also easy to add in temporal information using a time-delay embedding, and this is done for the same piece at the chordal level in Figure 11. Here the triplets of chords have some structure, with one circle when $0.25 < \epsilon < 0.32$ and three circles when $0.32 < \epsilon < 0.4$. It would be great to be able to interpret this kind of thing!

The final examples examine Bach's Chorale No. 19, with musical score shown in Figure 12.

A MIDI file of this piece, from [4], is parsed to extract the four voices. The dis-

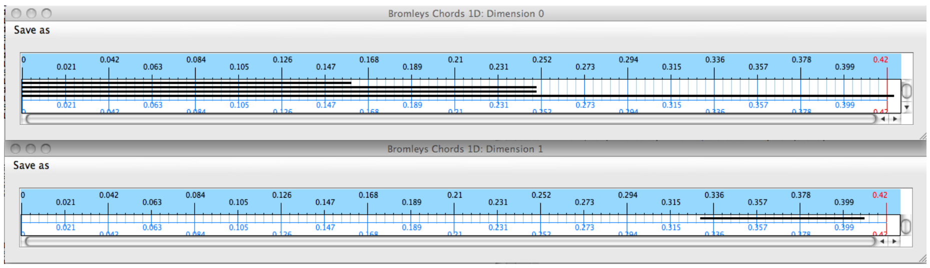

Figure 10: Barcodes for the chord progression from the score of "Abbott's Bromley Horn Dance" in Figure 6.

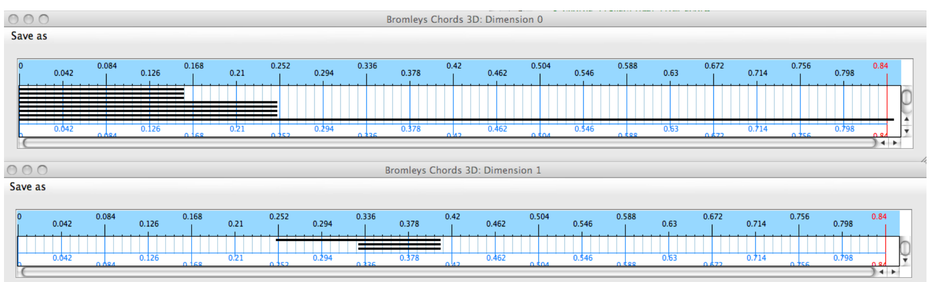

Figure 11: Barcodes for the dimension-three time-delay embedding of the chord progression from the score of "Abbott's Bromley Horn Dance."

Figure 12: The standard MIDI version of Bach's Chorale No. 19 is analyzed using the ideas of persistent homology in Figures 13 and 14.

tances between all four-part chords are calculated according to the chord-class metric (3), and the results are input to the Plex software in order to draw the barcodes. This is shown in Figure 13. The dimension 0 barcode shows a large number of chords that are separated by $\epsilon = 0.08$, a somewhat smaller number of chords that are separated by a distance of $\epsilon = 0.16$, and two chords separated by $\epsilon = 0.23$. Above this value, all chords merge into one connected component.

The dimension 1 barcode in Figure 13 shows $\text{Betti}_0 = 3$ connected components and one circle $\text{Betti}_1 = 1$ for $0.16 < \epsilon = 0.23$, and this structure then changes to $\text{Betti}_0 = 1$ and $\text{Betti}_1 = 3$ for $0.24 < \epsilon = 0.33$. Features such as these appear to be unique identifiers of the particular pieces, meaning that other Bach Chorales from the same data set have different Betti numbers that occur over different ranges of ϵ. Finding the origin of such variations is an interesting challenge.

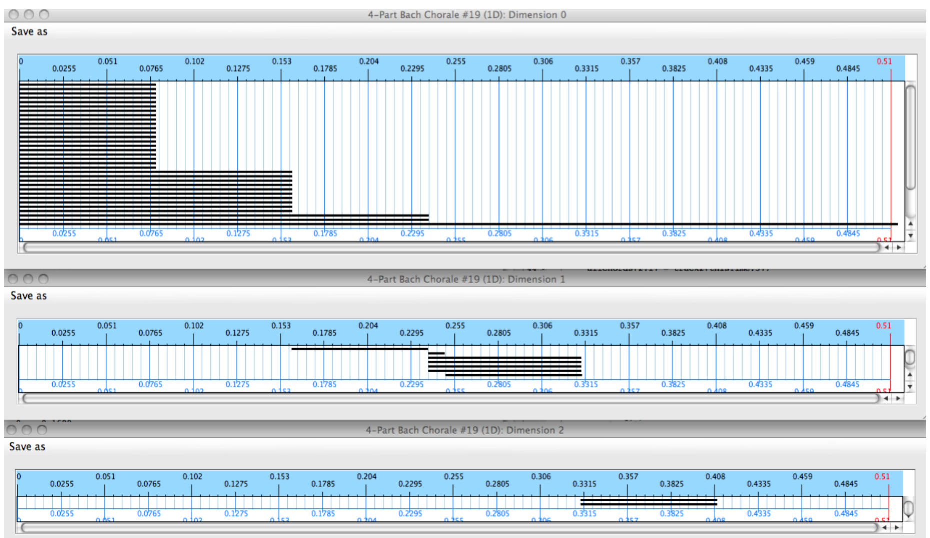

Figure 13: Barcodes for Bach's Chorale No. 19.

Finally, Figure 14 shows the two-dimensional time-delay embedding of the Bach Chorale, where the zero dimensional plot is interpretable directly in terms of the distribution of chord pairs and how they cluster under the chord-class metric. Again, there is a collection of persistent Betti_1 circles.

4 Rhythmic Barcodes and the Necklace Notation

Rhythmic notations represent time via a spatial metaphor. In standard musical notation, time is drawn linearly, though there are alternative notations such the necklace notation of Figure 3 that display the circular nature of rhythmic patterns: each pass through the cycle is one repetition of the rhythmic motif. The pitch-class metric (1)-(2) is not immediately applicable to the task of measuring such cycles because pitch is perceived

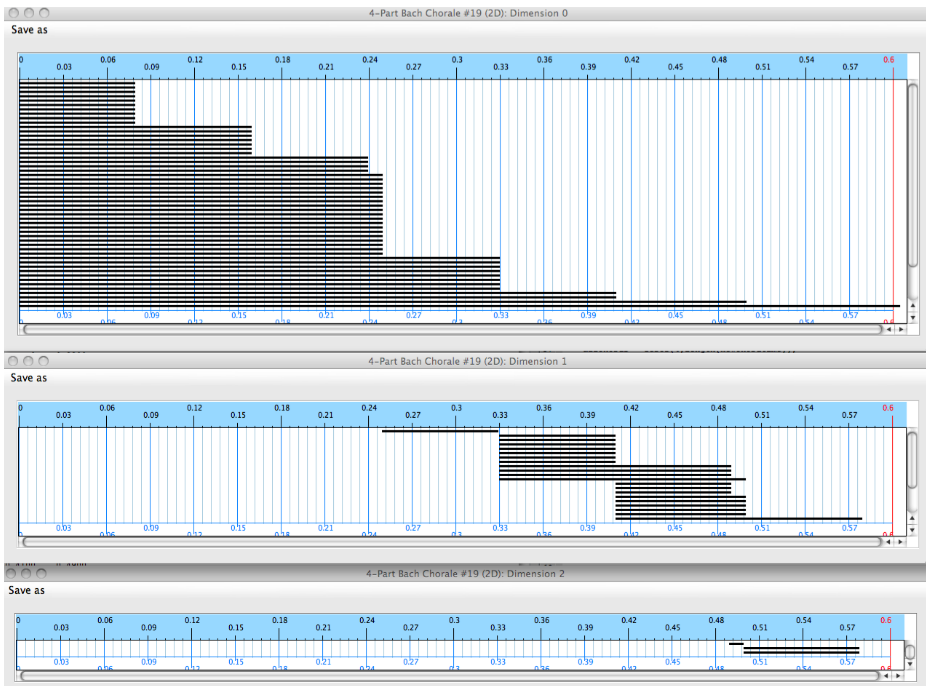

Figure 14: Barcodes for the two-dimensional time-delay embedding of Bach's Chorale No. 19.

in a logarithmic fashion while time is not. Accordingly, the metric can be modified to measure the distance between two times f and g as

$$d(f, g) = \min(s, 1 - s) \text{ where } s = \mod(|f - g|, 1), \tag{4}$$

and where one unit of time represents one period of the rhythm. Let's call this the *necklace metric*.

The "Ewe" rhythm of Figure 3 is translated into the vector of time points

$$\{0, \frac{1}{6}, \frac{2}{6}, \frac{5}{12}, \frac{7}{12}, \frac{3}{4}, \frac{11}{12}\}. \tag{5}$$

The distance is calculated between all of these time points under the necklace metric (4), and this set of distances is input into the Plex software. The resulting barcodes are shown in Figure 15.

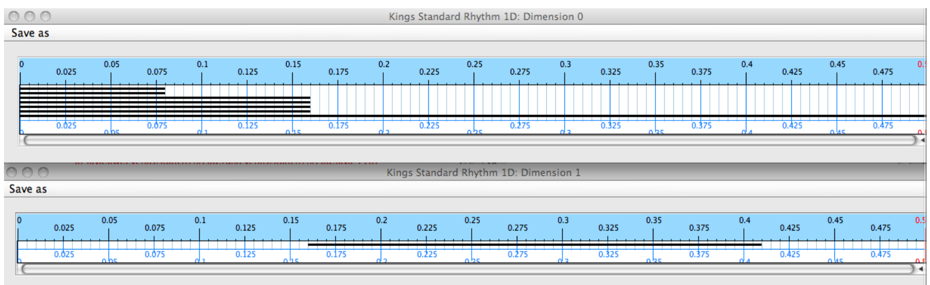

Figure 15: Barcodes for the "standard rhythm" [5] of Figure 3 show the distribution of time intervals in the rhythm in the top (dimension 0) plot and show the circular structure with $\text{Betti}_0 = 1$ and $\text{Betti}_1 = 1$ in the bottom (dimension 1) plot.

The dimension 0 barcode shows the clustering of the points in time. In the sequence (5), the minimum distance is $\frac{1}{12}$, and this occurs in two places, between the third and fourth notes, and again between the 11th and the first notes. Accordingly, the barcode shows two lines that vanish when ϵ reaches 0.08. Since the largest distance between any two adjacent time points is 0.16, all the points merge into one cluster at $\epsilon = 0.16$. The dimension 1 barcode displays a persistent Betti_1 bar from $0.16 < \epsilon < 0.42$. This is the anticipated cycle around the necklace.

As might be expected, more complex rhythmic patterns and higher dimensional embeddings yield more complex barcodes. For instance, the rhythm of the first 4 measures (a 24 beat cycle) of "Abbott's Bromley Horn Dance" are shown in Figure 16. As usual, the distribution of short and long intervals is shown in the dimension 0 barcode while the circular structure of the rhythm appears in the dimension 1 barcode for $0.08 < \epsilon < 0.32$. There are also interesting features to this rhythm in the two and three dimensional barcodes.

Figure 16: Barcodes for the first four measure of the traditional folk tune "Abbott's Bromley Horn Dance" of Figure 6 show the distribution of time intervals in the rhythm in the top (dimension 0) plot and show the circular structure with $\text{Betti}_0 = 1$ and $\text{Betti}_1 = 1$ in the second (dimension 1) plot. Higher dimensional features are also readily apparent.

5 Discussion

This paper argues that an investigation of the topological structures inherent in musical data is feasible using the ideas of persistent homology. Besides demonstrating that well known topological features can be derived from musical data sets, such analyses may be useful in information retrieval, in analysis of musical pieces, and in applications such as audio segmentation, melody recognition, and musical classification. Musical data may be ideally suited as a vehicle for exploration of the techniques of persistent homology because it is obvious (at least in retrospect) that there are significant topological structures present. Three examples are given: the circle of notes shown in Figure 1, the circle of fifths shown in Figure 2, and the circular form of the rhythmic necklace notation shown in Figure 3. These are readily identifiable in the barcodes derived from musical data displayed in Figures 4, 9, and 15.

There are a number of issues raised that may lead in fruitful directions. Generalizing the homological analysis to other musical domains such as spectral and audio rate data would provide an important and nontrivial extension. For instance, harmonic musical instruments (such as those that make sounds using strings or air columns) have waveforms that are approximately periodic. Since periodic waves can be pictured as a function on the circle, they should exhibit nontrivial topological structure. A basic question is what metrics can be applied in these alternative domains since different metrics may be able to provide different kinds of information. It is likely that the requirement that the distance be given between all pairs of points in the data set is overrestrictive. If this can be relaxed, it would allow the use of partial orderings, and might be appropriate for submajorization [15]. This could be especially useful for perceptual data that does not conform to a metric or where the information may be incomplete.

Even with barcodes at the symbolic (.mid) level, questions remain. It is not yet clear what kind of geometric shapes correspond to the higher level Betti numbers in the more complex barcodes (such as in Figures 7 and 8). From the musical perspective, it is important to ask how such topological structures can be interpreted in terms of the underlying musical piece. A good approach might be to evaluate a larger corpus of melodies, harmonies, and or rhythms with the goal of fully deciphering such relationships. Even if higher dimensional barcodes cannot be interpreted easily, they might be useful in classification as a kind of signature or "feature vector" for subsequent processing. Similarly, they might be useful in automatic segmentation to determine when something has changed (for instance in the underlying scale or the underlying melodic pattern).

There are other ways of incorporating temporal information than using the time-delay embedding, and these might also help to provide a fuller topological analysis that includes both pitch and rhythmic analyses simultaneously. Similarly, rhythmic patterns often occur in hierarchies and it would be interesting to pursue the idea of locating persistent homological structures from hierarchical musical data.

References

[1] Safî al-Din al-Urmawî, *Kitâb al-Adwâr* 1252, trans. R. Erlanger in *La Musique arabe*, Paul Geuthner, Paris, 1938.

[2] W. Anku, "Circles and time: a theory of structural organization of rhythm in African music," *Music Theory Online*, Vol. 6, No. 1, Jan. 2000.

[3] G. Carlsson, "Topology and Data," *Bulletin of the American Mathematical Society*, Volume 46, Number 2, April 2009, Pages 255-308.

[4] Classical MIDI Archives http://www.classicalarchives.com/

[5] A. King, "Employments of the "standard pattern" in Yoruba music." *African Music*, 2 (3) 1961.

[6] A.B. Lee, K.S. Pedersen, and D. Mumford, "The nonlinear statistics of high-contrast patches in natural images," *International Journal of Computer Vision* (54), No. 1-3, August 2003, pp. 83-103.

[7] C. Peterson, "Traditional Music," http://www.cpmusic.com/tradmus.html

[8] http://comptop.stanford.edu/programs/jplex/index.html

[9] John Rahn, *Basic Atonal Theory*, Prentice Hall International, 1980.

[10] http://en.wikipedia.org/wiki/Roman_numeral_analysis

[11] H. Schenker; ed. O. Jonas, *Harmony*, trans. E. Mann-Borgese. Chicago: University of Chicago Press, 1954.

[12] W.A. Sethares, *Tuning, Timbre, Spectrum, Scale*, Springer-Verlag, London, 2nd edn, 2004.

[13] W.A. Sethares, *Rhythm and Transforms*, Springer-Verlag, London, 2007.

[14] G. Singh, F. Memoli, T. Ishkhanov, G. Carlsson, G. Sapiro and D. Ringach, "Topological structure of population activity in primary visual cortex," *Journal of Vision*, Vol. 8, No. 8, (11), pp. 1-18, 2008.

[15] D. Tymoczko, *Geometry of Music*, 2010.

[16] http://en.wikipedia.org/wiki/Circle_of_fifths