Visual Hierarchical Key Analysis

CRAIG STUART SAPP
Center for Computer Assisted Research in the Humanities,
Center for Research in Music and Acoustics, Stanford University

Tonal music is often conceived of as progressing through a sequence of key regions, usually starting and ending in the tonic key, with a journey away from the tonic key somewhere in the middle of the piece. This article presents a visual method of displaying the musical key structure of a composition in a single picture. The hierarchical plots can also show the relative strength of these key regions and how they develop out of the chordal substrate of the music.

Categories and Subject Descriptors: H.5.5 [**Information Interfaces and Presentation**]: Sound and Music Computing

General Terms: Design, Algorithms, Experimentation, Human Factors, Theory

Additional Key Words and Phrases: Music visualization, tonality models, tonal analysis, harmonic analysis, computational music cognition

1. COMPUTATIONAL KEY IDENTIFICATION

Consider the melody shown in Figure 1:

Fig. 1. A short song (Liebes-A-B-C by August Pohlenz, 1827).

In what musical key is this melody? For a trained musician, it is easy to decide that the music is in the key of F major. This is due to the key signature containing one flat, the first note being an F, and the last note being an F as well. However, these are only superficial features used to determine the key of the music by eye. Suppose that someone listens to this melody without access to the score? Even in this case, musicians should easily be able to tell that the music is in a major key and that the first and last notes are the tonic of the key.

Now suppose that, in an attempt to hide the fact that the music is in F major, the music is altered by changing the key signature as well as the first and last notes, as demonstrated in Figure 2:

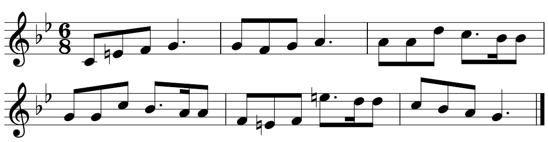

Fig. 2. The same melody as in Figure 1, but with a slightly altered beginning and ending.

Author's address: Center for Computer Assisted Research in the Humanities, Center for Computer Research in Music and Acoustics, Stanford University , Stanford, CA; Email: craig@ccrma.stanford.edu

Has the key changed? At a cursory glance, it might appear that the music is in G minor because the last note is G, the longest note in the first measure is G, and the key signature matches that of G minor. Only the first and the last few notes have been changed, while the key signature has just been changed to suggest G minor is the key of the music. Nonetheless, the key of the song has not changed. Listening to the last note of the song, a non-musician might perceive the end of the music as unresolved, and the last note might seem unstable, since it is expected to go down one more step. For a trained musician, the last note should sound like the second degree of a major scale, which is an unusual note on which to end a song.

So what is a more robust method of determining the key of a song? One method could be to look at interval patterns that are more likely to occur in one key rather than another. For example, the pattern F-E-F-G that starts the original version of the song can only occur in the two major keys of C and F. Another way of looking at this note sequence is by intervals: the sequence first goes down one half-step, then up one half-step and then up a whole-step. This can only occur in two places in a major key: on the scale-degree pattern 1-7-1-2 or the pattern 4-3-4-5. Taken in isolation, the scale-degree pattern 1-7-1-2 is much more common in music, so a musician should quickly suspect that the first measure of the original version of the song is in the key of F major (with the scale-degree pattern 1-7-1-2) rather than in C major (with the scale-degree pattern 4-3-4-5). This is a good method for determining the perception of a musical key in a melody; however, it is difficult to apply to more complicated music that contains harmonic accompaniment along with the melody, since the music might be more difficult to break down into linear sequences of interval patterns.

Another key-finding method, which works well for polyphonic music, was developed by Carol Krumhansl and Mark Schmuckler in the 1980s [Krumhansl 1990, ch. 4]. In their key-finding algorithm, the relative amount of pitches in the music is compared to a prototypical pattern of pitches expected in major and minor keys. For example, there obviously should be more pitches in a song that are scale-degree notes of the key rather than non-scale notes. If a piece of music is in C major, one would expect lots of pitches in the scale of C major (C, D, E, F, G, A, and B) and not expect many notes outside of the key: (C#, D#, F#, G#, and A#).

With knowledge about the number of notes in each pitch-class in a song, a key can be determined by *correlation* with a prototypical pattern for a major (or minor) key. Correlation is a mathematical means of determining how similar two patterns are to each other. Figure 3 shows a sequential pattern which will be used to find example correlations:

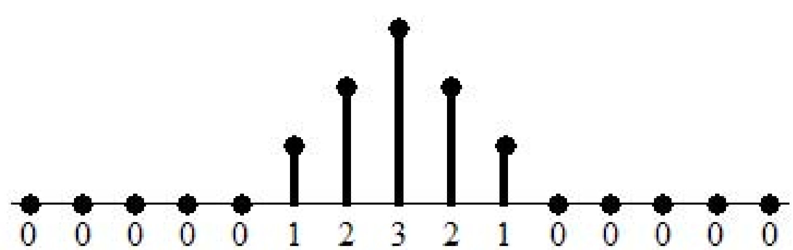

Fig. 3. A sample pattern sequence to be matched.

How similar is the pattern in Figure 3 to the following pattern in Figure 4?

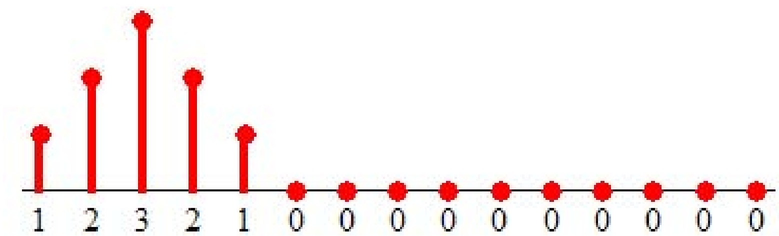

Fig. 4. A sequence to be compared with Figure 3.

To determine the correlation between these two patterns, multiply each corresponding element in each sequence, and then add the resulting values of the sequence together. Example calculations for the correlation between Figure 3 and Figure 4 are summarized below:

```
  0   0   0   0   0   1   2   3   2   1   0   0   0   0   0
× 1   2   3   2   1   0   0   0   0   0   0   0   0   0   0
 ___________________________________________________________
  0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0  = 0
```

Since the calculations yield a value of zero, there is no correlation between these two patterns. This might seem strange, since to the eye, they do seem similar as both have the same basic shape. However, the correlation value does not consider shifting the patterns, and is only looking point-by-point between the two patterns when comparing them. Now suppose that the second pattern is shifted slightly to the right so that there is a little bit of overlap between the patterns, in this case the correlation is a little larger:

```
  0   0   0   0   0   1   2   3   2   1   0   0   0   0   0
× 0   1   2   3   2   1   0   0   0   0   0   0   0   0   0
 ___________________________________________________________
  0 + 0 + 0 + 0 + 0 + 1 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0  = 1
```

In this case there is slightly more correlation (similarity) between the two patterns, but not very much. At what point is the correlation between the original pattern and the shifted pattern the greatest? Obviously when the peak of the shifted pattern is aligned with the first pattern:

```
  0   0   0   0   0   1   2   3   2   1   0   0   0   0   0
× 0   0   0   0   0   1   2   3   2   1   0   0   0   0   0
 ___________________________________________________________
  0 + 0 + 0 + 0 + 0 + 1 + 4 + 9 + 4 + 1 + 0 + 0 + 0 + 0 + 0  = 19
```

For our purposes in determining the musical key of a song, it is useful to plot all of the correlation values for each shift of the matching pattern. For example, Figure 5 displays a plot of the 11 correlation values generated by shifting the second pattern (Figure 4) gradually from the left against the original pattern (Figure 3).

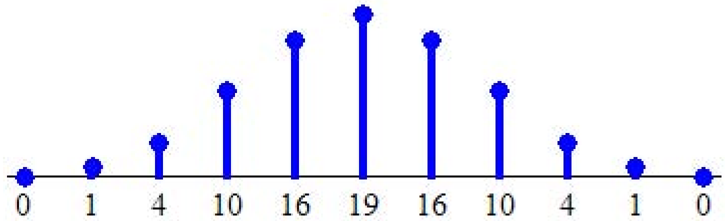

Fig. 5. Correlation values for shifted version of Figure 4 when compared to Figure 3.

How do correlation measurements apply when finding the key of a song? First, count all of the pitch classes (pitch names without octave information) to generate a pitch-class histogram. For example, when counting by sixteenth-note durations, there are 7 duration units of C in the original song from Figure 1; 0 duration units of C#; 5 duration units of D in the song, *etc*., as shown in the histogram in Figure 6.

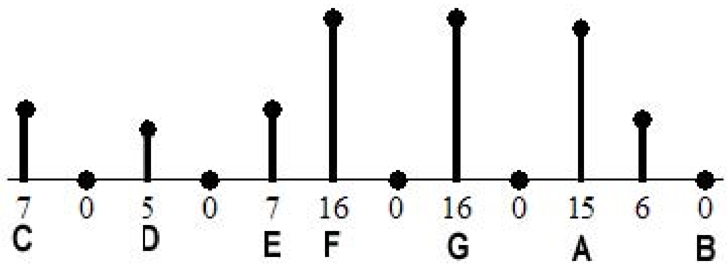

Fig. 6. Pitch-duration histogram for the music from Figure 1 in sixteenth-note durations.

This histogram of the pitch-classes in the original song (weighted by their duration) can now be correlated with the patterns typical for those in a major key. What are those values? There are many ways to derive a distribution of pitches for a major key, but the basic Krumansl-Schmuckler key-finding algorithm in particular uses probe-tone weights from perceptual experiments as the major scale pattern, shown in Figure 7.

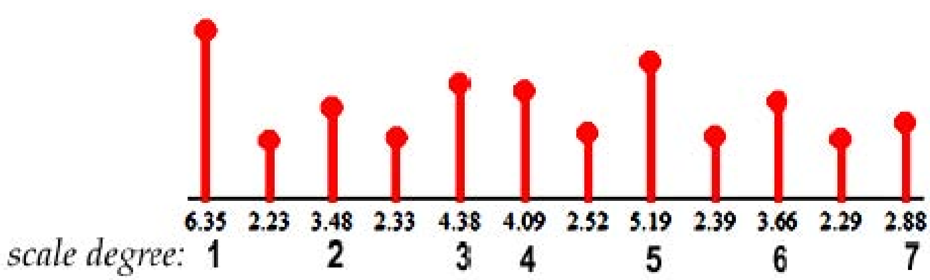

Fig. 7. Krumhansl's probe-tone weightings for major key contexts.

Note that the scale-degree positions have higher weighting values than do the non-scale-degree values. If the same correlation calculations are now done as in Figure 5 (but wrapping values that extend to the left back onto the right side of the sequence), the correlation values for each shifted version of the major pattern will be generated as shown in Figure 8.

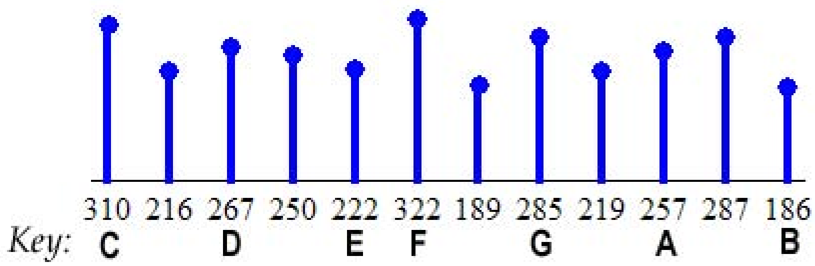

Fig. 8. Correlation values for rotated versions of an expected major scale pitch distribution (Figure 7) when compared to the histogram of pitch-classes from a real sample of music (Figure 6).

Since the highest correlation value (322) is generated when the song's pitch histogram and the major-key profile occurs when the tonic note is on F, the musical key that best fits the music is F major.

The previous discussion demonstrates the essence of key-finding by correlation. The mathematical description of the correlation calculations being done for the above example key-finding calculations is

$$\text{key} = \max_{j=0}^{N-1} \sum_{i} e(i,j)d(i)$$

where *e* is the major (or minor) key profile (such as the one in Figure 7); *d* is the histogram of pitches in the actual music (such as that in Figure 6); *i* is the set of pitch-classes; *j* is the shifting of the scale profile onto the various pitch-classes; and *N* is 12, the number of pitch-classes. The full practical implementation of the Krumhansl-Schmuckler algorithm normalizes the correlation values into the range between +1.0, and -1.0, where +1.0 means maximally correlated, and -1.0 maximally uncorrelated, using the normalizations in the following normalized correlation equation:

$$\text{key} = \max_{j=0}^{N-1} \frac{\sum_{i=0}^{N-1}[e(i,j)-\mu_e][d(i)-\mu_d]}{\sqrt{\sum_{i=0}^{N-1}[e(i,j)-\mu_e]^2 \sum_{i=0}^{N-1}[d(i)-\mu_d]^2}}$$

where μ_e is the arithmetic mean (average) of the *e* measured pitch profile, and μ_d is the arithmetic mean of the *d* histogram. In the field of statistics the values calculated with this normalization are called r-values. The r-value is a correlation coefficient that indicates the strength of a relationship between data sets (pitch histograms and key profiles in our case). An r-value of +1.0 means that the data sets are very strongly related; an r-value of 0 means that there is no relation between data sets. The r-value normalization of the correlation allows comparisons between different key patterns, such as the probe-tone profile for minor keys in the Krumhansl-Schmuckler key-finding algorithm (Figure 9).

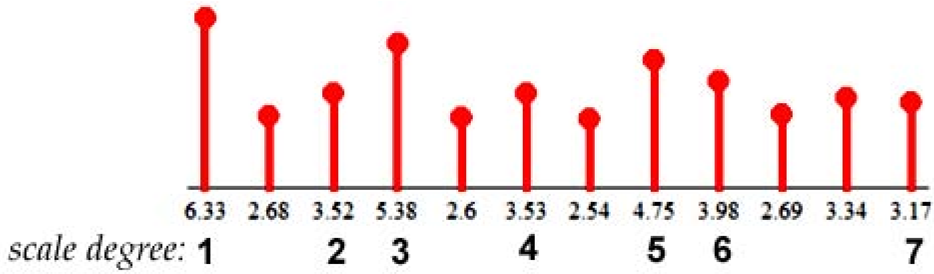

Fig. 9. Krumhansl's probe-tone weightings for minor-key contexts.

2. MUSICAL KEY MODULATION

Correlation is a good technique for finding the key in a selection of music that entirely or predominantly remains in the same key, but when more that one key is present in the music, the meaning of the correlation between a pitch histogram and a prototypical key profile is less reliable and can be difficult to interpret. The correlation technique enumerates the possible keys in the musical selection, but it cannot by itself identify cases when there are supposed to be two or more keys present in a musical sample.

Take, for example, the sixth variation from Franz Schubert's 13 Variations on a Theme by Anselm Hüttenbrenner in A major, D. 576 (1817) which starts out in the key of F# minor and ends in the key of A major (Figure 10). Even if you are not a trained musician, you should still be able to hear the character of the music change half-way through the variation as the music transitions from a sadder, melancholy feeling into a happier, optimistic feeling half-way through the variation.

Fig. 10. Franz Schubert’s 13 Variations on a Theme by Anselm Hütenbrenner in A minor, D. 576 (1817), var. no. 6.

If a pitch histogram of the entire variation is given as input to the Krumhansl-Schmuckler key-finding algorithm, it will determine that the best key for the music is in A major, with an r-value of 0.86. F# minor (r=0.78) is the second-best choice, followed by C# minor (r=0.62), and E major (r=0.57). The key-finding algorithm identifies A major as the best key for the music, but it cannot indicate the internal key structure of the music except by giving F# minor, which has equal importance to A major in the variation, as the second best key to fit the entire variation.

Now suppose that the music of the variation is split into two halves, so that key of each half can be examined separately via the Krumhansl-Schmuckler key-finding algorithm. In this case the algorithm will find the correct key assignments for each half. The first half will be identified as F# minor (r=0.80); the second half will be identified as A major (r=0.90).

If the music is further split into a sequence of four four-measure segments, the Krumhansl-Schmuckler key-finding algorithm will identify the key sequence: C# minor (r=0.67), F# minor (r=0.87), E major (r=0.77), A major (r=0.84). This key analysis does not make precise musical sense, but it does describe the basic movement from dominant to tonic key regions in each section (C# to F#) and (E to A). In this case, not enough information is present in the pitch histograms to give the same answer for the first and second halves of each key region, so the Krumhansl-Schmuckler algorithm is looking at a slightly lower level of music resolution than the actual key, and it displays the overall chord structure from dominant to tonic for each key.

In the first case, where all of the music for the variation was used to estimate the key, only one key in the music was identified because the Krumhansl-Schmuckler analysis window was too large and included two key regions. The algorithm chose A major, perhaps only because there are more note-heads in the second half of the piece, or because of the algorithm's tendency to emphasize mediant key areas more than sub-mediant key areas when using the standard probe-tone weightings. In the second case, the algorithm correctly identified the two keys in the music when the music was cut into two equal parts, exactly proscribing the two key areas in the variation. In the third case, the duration of music input into the Krumhansl-Schmuckler algorithm calculations was too small, and a sub-key description of the music was given as output.

Note that if too much music is analyzed at once, fewer important keys are suppressed; if too little music is analyzed at once, the chordal structure of the music is really being analyzed instead of the key structure. So what is the proper duration of music to give as input to the Krumhansl-Schmuckler key-finding algorithm?

3. KEY ANALYSIS PLOTS

This author's approach is to examine *all possible segmentations* of the music. To accomplish this, a two-dimensional plot was devised to display the analysis results from thousands of key identifications generated by the Krumhansl-Schmuckler key-finding algorithm, or any other computational key-finding algorithm.

Figure 11 shows the layout of the plot domain used to display the key analyses from a computational key-finding algorithm. The horizontal axis represents time in the music, from the start of the piece at the far left side through the end of the piece on the far right side. The vertical axis represents the duration of music given to the key-finding algorithm. For example, the top of the plot represent the entire duration of the composition, while the bottom of the plot represents some minimum time duration, such as one beat in the music. The duration of music represented on the vertical axis is also called the *analysis windows size*.

To generate data for plot, an analysis window is selected on the vertical axis and centered at a time in the composition represented on the horizontal axis. For example, the point in Figure 11 labeled *first half of composition*, represents the key analysis of the first half of the composition. The key analysis for that window is displayed in the plot in the middle of the analysis window's duration in the music. For the case of the Schubert variation, the first half of the music is found to be in F# minor, so that point in the plot will be labeled F# minor. Likewise, the point marked *second half of composition* will be labeled as being in A major. To get the intervening points, an analysis window equal to 1/2 of the composition gradually slides through the music from start to finish; an analysis

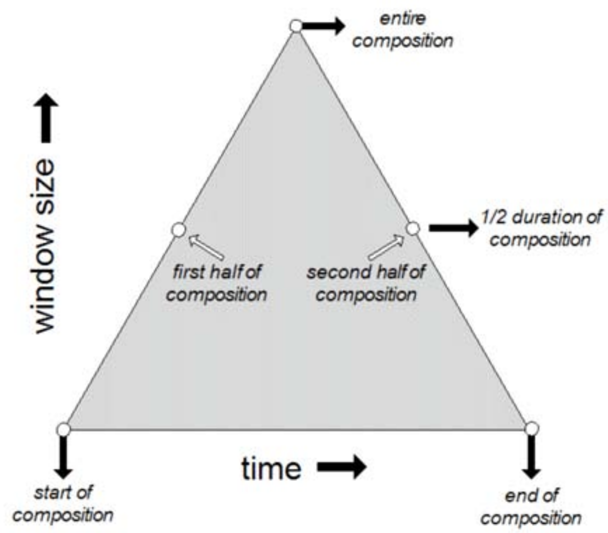

Fig. 11. The significance of the horizontal and vertical axes in a key-analysis plot with interpretations of a few points in the plot.

result is labeled in the plot at the mid-point time location of the window in the music on the horizontal axis.

Every point in the plotting region represents a key analysis done with two parameters: (1) the duration of the analysis window into the music; and (2) the center-point in time of the analysis window. The plotting region is in the shape of a triangle because each analysis-window size represented on the vertical axis is centered at a specific time in the composition. For example, the top of the plot can only have one point, since there is only one way of centering an analysis window of the entire piece (at the half-way point in the composition).

Since every point in the plotting domain shown in Figure 11 is a separate key analysis, it would be difficult to display analysis results in textual form. Instead, each key analysis result is assigned a color. What color to assign each of the 24 major/minor keys is an arbitrary decision. For the following plots, the coloring convention given in Figure 12 is used, where the diatonic circle of fifths is mapped onto the colors of the rainbow. This allows harmonic relations of a fifth, which are the most important relations in tonal harmony, to be represented by closely related colors. Chromatic alterations of the diatonic pitches are given as slight variations in the color of the natural pitch-classes, as indicated by the accidentals below the diatonic E colors in Figure 12. Major and minor keys can be distinguished by brightness as also illustrated in Figure 12.

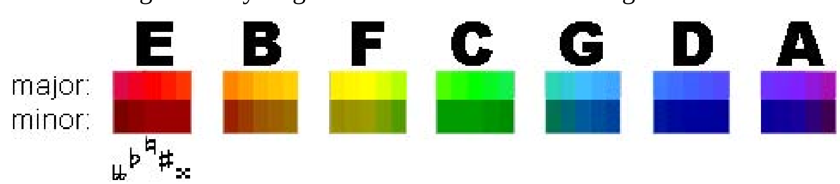

Fig. 12. Key-to-color mappings used in this paper. Diatonic pitches are arranged in a sequence of fifths on the colors of the rainbow. Chromatic alterations of diatonic pitches generate slight color shifts, and modality is indicated by brightness.

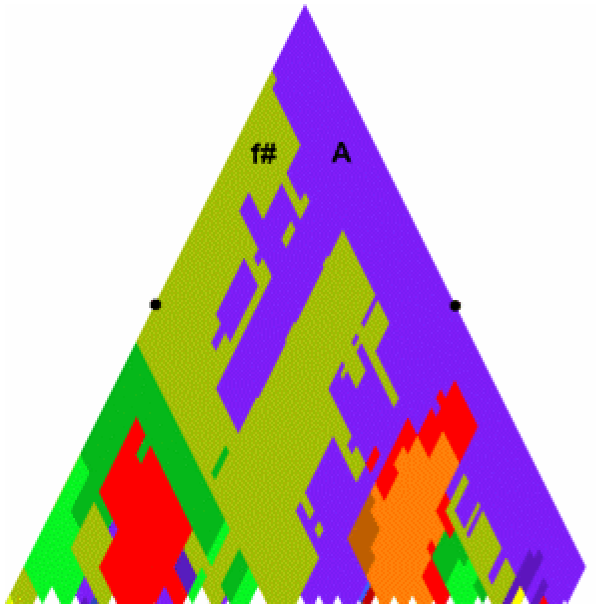

Fig. 13. Linear keyscape plot of the music shown in Figure 10. The key regions of F# minor and A major are labeled with their tonic letters. The black dots in the plot indicate the analysis window for the first half of the music (left side) which is F# minor, and the last half of the music (right side) which is in A major.

Using the coloring scheme in Figure 12, key analyses of the music in Figure 10 are displayed in the plot shown in Figure 13. Note that the two black dots in Figure 13 indicate the first and second halves of the variation. The Krumhansl-Schmuckler algorithm correctly identifies the two keys of the variation at these points. The first half of the music is in F# minor, which in the plot is colored in a dark greenish-yellow color. The second half of the music is in A major, colored in purple. Also notice that the colors of the plot in the immediate vicinity of these two points also agree with the key-analyses at the two given points, which shows there is some tolerance for the presence of non-key materials in the analysis window in the key-finding algorithm. Other colors in the plot are due to temporary key regions and/or chord roots, since the bottom of the plots represent small-scale key features in the music.

Figure 14 shows the same plot as Figure 13 and illustrates the input data used for the key-finding algorithm at various points in the plot. These types of plots are given the nickname *keyscapes,* since they are analogous to landscape paintings. The bottom of the plots represent small-scale key features such as chords that are similar to the foreground in a landscape painting. The top of each plot represents the large-scale key of the composition, which is similar to the background in a landscape painting. The various vertical positions in the plots are also related to the concept of foreground, middleground, and background in Schenkerian analysis, which is a form of tonal music analysis developed in the early 20th century. Musical keyscape plots, to some extent, serve as an objective form of Schenkerian analysis [see Narmour 1977].

The keyscape plots were originally conceived for comparing the properties of various key-finding algorithms at various time resolutions. For example, Aarden [2003] has measured a different set of major/minor key correlation profiles from musical scores, which can be used in place of the Krumhansl probe-tone profiles in the correlation algorithm. Figure 15 shows Aarden’s major and minor profiles used to find the key by correlation calculations. Using the major/minor profiles from Figure 15, a keyscape plot

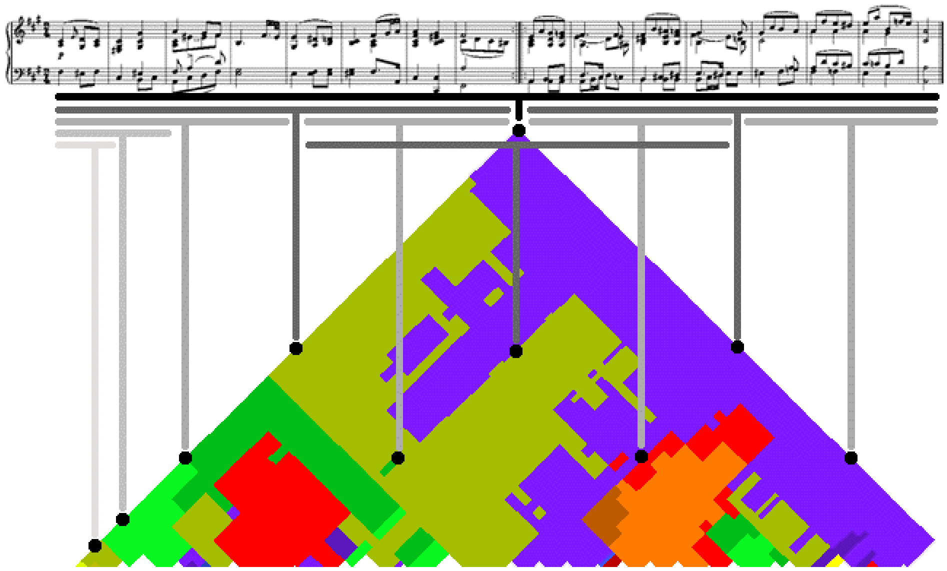

Fig. 14. Keyscape plot of the music from Figure 13 and the relationship of various points in the plot to music from Figure 10.

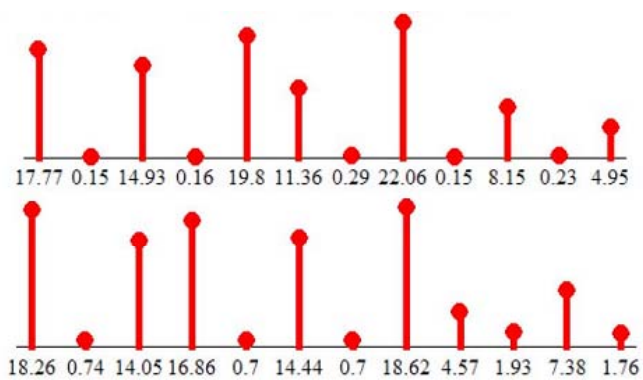

Fig. 15. Aarden’s correlation profiles for major (top) and minor (bottom) keys.

for the Aarden profiles is shown in Figure 16. In this plot the F# minor half of the music is identified as stronger than the A major half, since the top of the plot is colored for F# minor. The two primary key regions do have the same correct analyses of F# minor and A major in both plots at the appropriately marked points, but in the presence of modulation between keys, both differ in their analyses due to their different correlation weights.

Included with this paper is an animation graphic that interpolates the key profile weights between those of Krumhansl and Aarden. The regions of the plot that remain in a single color (and thus in a single key) can be interpreted as being more stable. Regions of the plot that shift colors during the morphing between the two profiles are less certain to be given the correct key assignment by either set of weights, and could even indicate the presence of a modulation boundary between adjacent key regions.

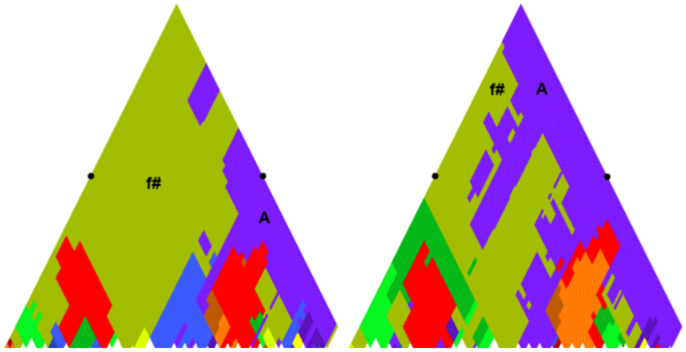

Fig. 16. Plot generated from Aarden weightings (left) compared to Krumhansl weightings (right) from Figure 13. Black dots in each plot indicate the two key regions of F# minor and A major in the music

Any key-finding algorithm can be input into the plot, not just those based on key-profile correlations. For example, Figure 17 plots the output from a root-finding algorithm that is independent of key-profiles. This algorithm also gives the correct key answers for the first and last halves of the variation. The Krumhansl-Schmuckler key-finding algorithm is reasonably accurate at finding the key despite the amazing fact that all notes in the analysis windows are jumbled together into a single histogram, just as if all the notes were played at once in a single large chord. Another key-finding algorithm that is sensitive to the ordering of the notes and chords was developed by Temperley [2002], and will be input into the keyscape plotting format in the future.

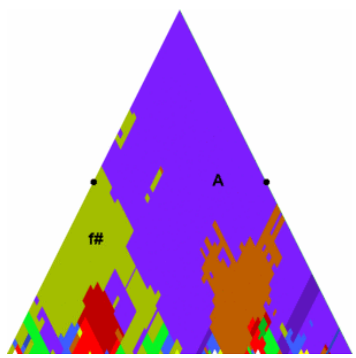

Fig. 17. Plot of the same music as in Figure 16, but using an unrelated root-identification algorithm.

Other types of descriptive information from music can be displayed in the keyscape plots. For example, ambiguity and clarity measurements derived from computational key-analysis algorithms can be used [Sapp 2001]. And the basic plotting domain from the keyscape pictures has been used to analyze musical features other than the keys [Segnini 2005]

4. TONAL STRUCTURE IN MUSIC

Most pieces of western classical music do not remain in one key throughout the entire composition. The typical format is to start in the tonic, go off into other keys, and then return to the initial key at the end of the piece. Older musical styles (say before 1800) are

Fig. 18. W.A. Mozart's Divertimento No. 4, K 439b (1783), mvmt. 1.

typically more conservative in the range of keys covered in a single composition, while later styles of classical music tend to get as far away from the initial key as possible (particularly in the music of the late 19th and early 20th centuries). With the keyscape plots, it is possible to get an overview of the harmonic structure of a piece of music and a rough estimate of the musical style in which the music was composed.

Figure 18 shows the musical example used in this section: Divertimento no. 4, K 439b, mvmt. 1 composed by W.A. Mozart. The form of this piece is an abbreviated sonata form. The first half of the music contains the first theme in C major (the tonic key) and the second theme in G major (the dominant key). The second half of the music starts with a development section that is intentionally ambiguous in terms of key. The development section starts with a phrase that is not really in any particular key but can be assigned to D minor for convenience. This is followed by two phrases in F major. The recapitulation section follows the development and repeats the second theme from the first half of the music; however the theme has been transposed from the dominant key into the tonic, as is traditionally done in the sonata form.

The music can best be described as going through the key sequence C-G-(d)-F-C. Summarizing the relative importance of the key regions is usually not done carefully in traditional music theory, but one simple way of describing the importance of each key is to indicate the number of measures that each key region lasts: C[15]-G[13]-d[4]-F[7]-C[24]. This gives a rough overview of the relative importance of each key region; however, this is similar to giving a table of numbers rather than a plot of numbers. The keyscape plots can be used to show the relative importance of key regions in a composition. Figure 19 displays a keyscape plot of the Mozart divertimento. Notice that the predominant color in the plot is bright green, which represents C major, the tonic key in the divertimento. The green section of the plot goes from the beginning of the music on the far left through to the end of the music on the far right, which indicates that the music starts and ends in C major. The very top of the plot is also colored green, which indicates that the key-finding algorithm determined that the entire composition as a whole was in C major.

The second most important key in the music is G major which is the key of the second theme in the introduction. This key is represented by light blue in the plot. Light blue is also the second most prominent color in the plot, agreeing with the music analysis of the

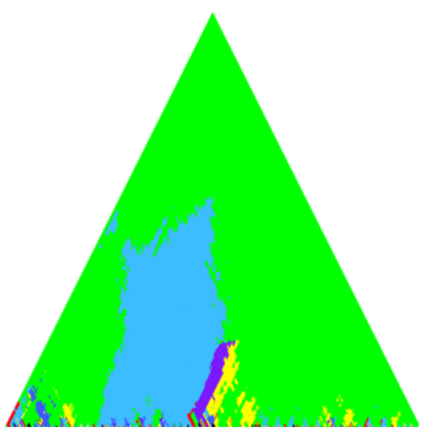

Fig. 19. Linear keyscape plot of the music found in Figure 18.

composition. The keys of the development are less important. This is shown by the violet and yellow sections on the lower right side of the light blue region in the plot. These color regions are smaller due to the key regions in the development section being shorter and less important than the key regions in the introduction. Notice that the development key regions do not go as high in the plot as the G major (light blue) region, since they are not as prominent in the music. Likewise, G major is not as prominent as C major represented by the green region, which reaches all the way to the top of the plot. The various colored regions at the bottom of the plot represent individual chord roots that merge into the weaker and then stronger key regions in the music. Thus the keyscape plots display the harmonic structure of the music in a hierarchical manner that is similar to Tree-Notation reductions of music found in *A Generative Theory of Tonal Music* [Lerdahl 1983].

Up to this point, the keyscape plots have been presented in a triangular format, but it is sometimes useful to display them with a rounded top, as found in Figure 20. Each row in a triangular form of the plot represents a window size that increases/decreases at a constant arithmetic rate. For example, if there are one thousand rows in a plot and the plot represents a piece of music with one thousand beats, the analysis window size for each row going upwards in the plot gets one beat longer on each row. Notice that the relative size of the window grows at a slower rate for larger windows. For example, the top row is about 0.1% larger than the row underneath it (1000 beats compared to 999 beats); however, the bottom row is 50% smaller than the row above it (1 beat compared to 2 beats). The plot in Figure 20 is rounded on the top because the size of each row increases at the same rate in a geometric progression rather than in an arithmetic progression. Therefore, the vertical scale of the plot in Figure 20 is logarithmic with respect to the size of the analysis window.

Since all levels of the music are weighted equally by the logarithmic scaling in Figure 20, it gives a more perceptually accurate view of the harmonic structure of the music. The triangle pictures are useful for viewing the large-scale key structure of a piece, or for short pieces such as the Schubert variation where the small-scale structure is also visible due to the short duration of the music. Figure 21 shows the Krumhansl and Aarden weighted plots of the divertimento movement. Looking at these plots, the Krumhansl weights can be seen to over-emphasize the dominant key (G major; light blue) at the expense of the sub-dominant key (F major; yellow). The Aarden weightings come closer to the music-theory expectations for the movement, since the sub-dominant key area in the development (yellow regions) is more prominent; however, the dominant region suffers a bit since it is only about as large as the briefer sub-dominant key region.

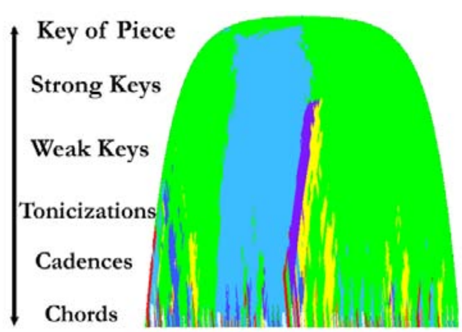

Fig. 20. Mozart divertimento with a logarithmic scaling of the analysis window size on the vertical axis.

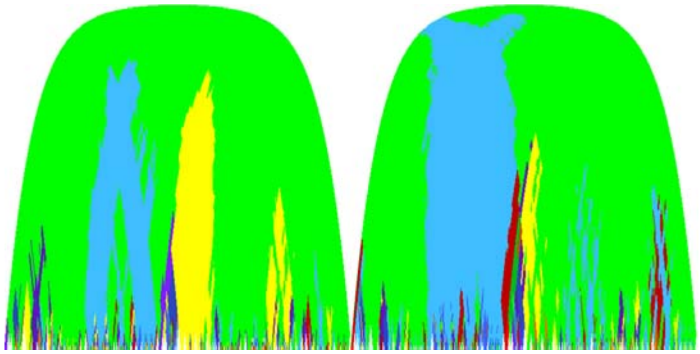

Fig. 21. Mozart divertimento with a logarithmic vertical scale comparing Aarden weights on the left and Krumhansl weights on the right.

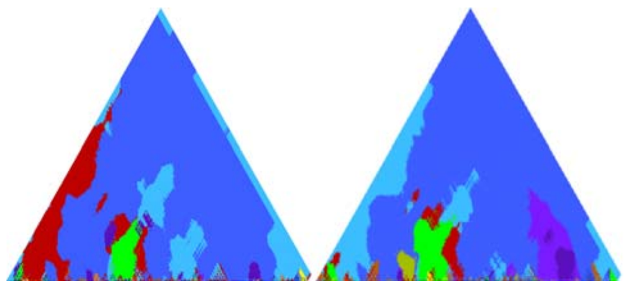

Fig. 22. Keyscape plots from the prelude from J.S. Bach's cello suite, BWV 1007. Aarden weights used on the left and Krumhansl weights on the right.

5. EXAMPLES FROM VARIOUS STYLES OF MUSIC

5.1 Analysis of the Prelude from J.S. Bach's Cello Suite, BWV 1007

Figure 22 shows the keyscape analysis of the prelude from J.S. Bach's cello suite, BWV 1007, which was composed around 1720. Both keyscape plots agree that the music starts and ends in the key of G major (light blue) with a significant region of the dominant (medium blue) in the middle of the piece. The Aarden weightings over-emphasize the sub-mediant chord (red), which is probably not that noticeable to human listeners.

The Aarden-weight plot in Figure 22 has a nice feature: at the top of the plot there is a small region of light blue which represents the key of G major. This means that when the entire composition is given to the algorithm, it identifies the key as G major, even though more dominant (medium blue) chords are present in the music. On the other hand, the Krumhansl weightings will identify D major as the key of the piece when given the notes of the entire movement, most likely due to an over-emphasis of the dominant key region.

The prelude has a fair number of non-harmonic tones, and since it is for a solo instrument, the music is sparse. Figure 23 shows comparable keyscape plots, with the music reduced to the basic chord sequence from each measure, which removes non-harmonic notes from the analysis. Both plots remain similar to the plots for the full version of the music, although the Krumahsl-weight plot does give more emphasis to the sub-median (red).

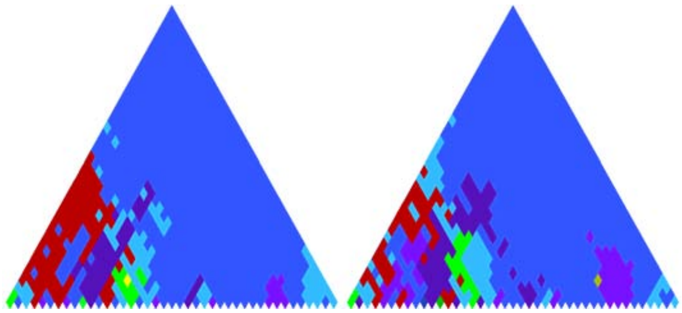

Fig. 23. Keyscape plots for a chordal reduction of the prelude also used in Figure 22. Aarden weights used on the left and Krumhansl weights on the right.

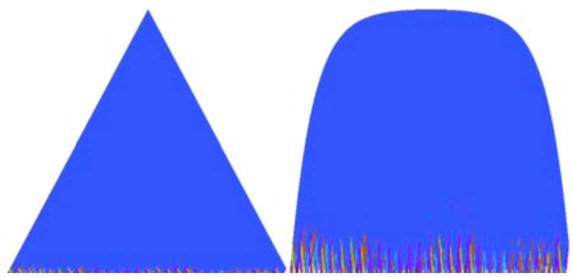

Fig. 24. Keyscape plots for Pachelbel's Canon in D Major in both linear (left) and logarithmic (right) vertical scaling, using Krumhansl weightings.

5.2 Analysis of Johann Pachelbel's Canon in D Major

By examining the plots in Figure 24 we can see that the music to Pachelbel's Canon is indeed in the key of D major, the key represented by the medium-blue color. According to the plots, there are absolutely no secondary key regions in the piece, since the key of D major arises directly from the underlying chords which are continually repeated.

In this case it is more interesting to display a plot of the music for every two bars (Figure 25), since the canon is a round that cycles every two bars. Each repetition has the same harmonic structure based on the baseline D-A-B-F#-G-D-G-A. There is some variation in the harmonic structure, which is made evident by the greenish-yellow regions representing a chord of F# minor in the plots of Figure 25. Sometimes, the greenish-yellow regions are not present in a segment, such as the 9th through 14th repetitions on the second row of the figure. In these cases the chord with a root on F# is changed into a chord with a root on D.

5.3 Analysis of Samuel Barber's Adagio for Strings

Figure 26 shows the keyscape plots for Samuel Barber's Adagio for Strings, written in 1936. The composition is in B-flat minor, which is indicated in the plots by the orange color extending to the top of the plots. The climax of the piece is visible as a red rectangle at the bottom of the plot near the end of the piece. This rectangle represents the key of F-flat major, although red is usually the color for the E pitch-classes (E and F-flat are the same key on a piano). Interestingly, F-flat is the most distant key from B-flat, since the

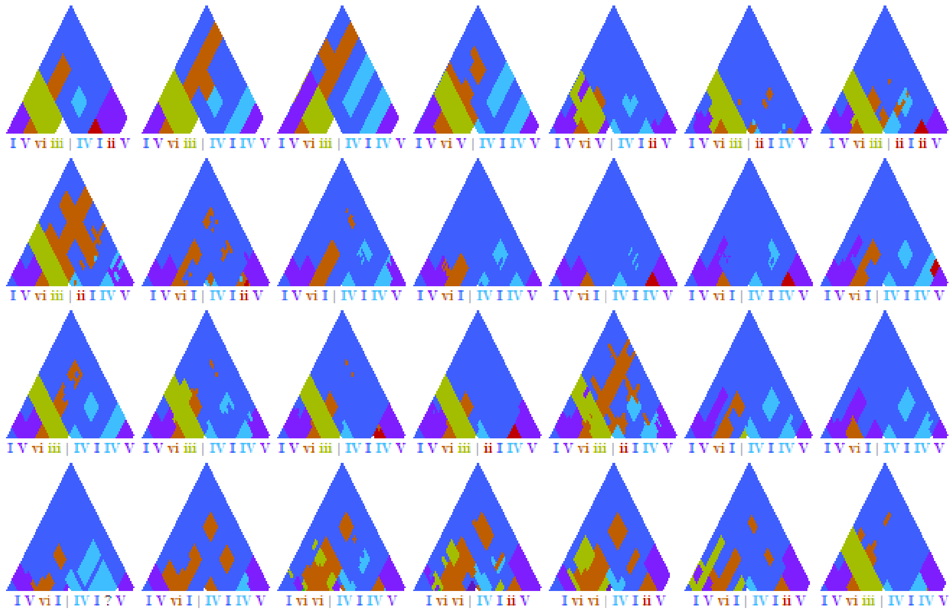

Fig. 25. Keyscape plots for every two bars of Pachelbel's Canon. Harmonic analyses are displayed underneath each plot, colored according to the chord root.

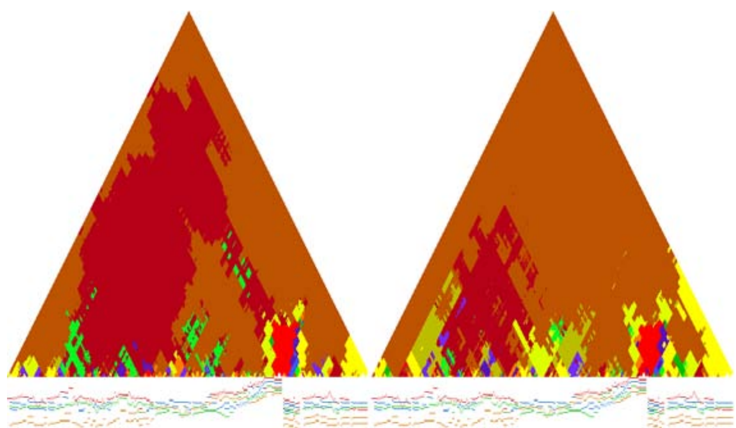

Fig. 26. Keyscape plots for Samuel Barber's Adagio for Strings. Aarden weights used on the left and Krumhansl weights on the right. Underneath each plot is a piano-roll notation of the music, where the vertical axis represents pitch from low to high.

interval between them is a tri-tone, but F-flat is played closely after the clearest presentation of a B-flat minor chord in the piece (measures 50-52).

The harmonic climax coincides with a large registral shift in the music. The lines below each plot in Figure 26 are plots of the notes present in the score. The horizontal axis represents time, just as in the keyscape plots. The vertical axis represents pitch, from low to high, and color is used to indicate the different instrumental parts. Notice that starting around the middle of the composition, the parts all start to rise gradually in pitch until they suddenly drop back to the starting register at about 75% of the way through the piece at the same time as the harmonic climax in the music.

Harmonically, the Adagio for Strings is an unusual composition. While the music is in B-flat minor, and both keyscape plots in Figure 26 agree with this analysis, the tonic chord of b-flat minor is used much less often than expected in a tonal piece. The opening of the piece starts with a single B-flat note, followed by a iv7 chord and then a V chord, giving a weak introduction to the key of the piece. Another unusual tonal device is ending the piece in the dominant instead of the tonic, which can be seen in the yellow colored region on the bottom right of the plots.

In the piece, the tonic chord is avoided as much as possible. For example, during the entire piece, the full tonic chord is only played 19 of the 564 quarter-note durations in the piece (3.3% of the total duration). On the other hand, the sub-dominant (dark red in the plots) and the dominant (yellow) are each played about three times more often than the tonic chord, as demonstrated in Figure 27. In a typical tonal composition, more emphasis on the tonic chord and less on the sub-dominant chord would be expected.

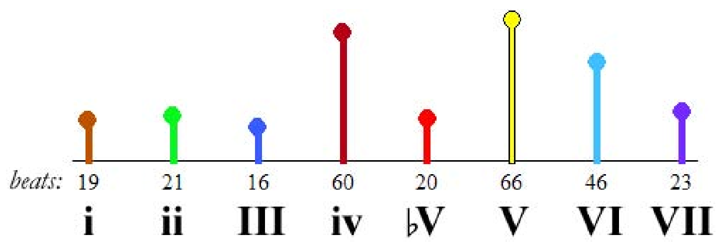

Fig. 27. Relative durations of prominent chord sonorities present in the Adagio for Strings. Colors represents chord roots which match the key colors in Figure 26.

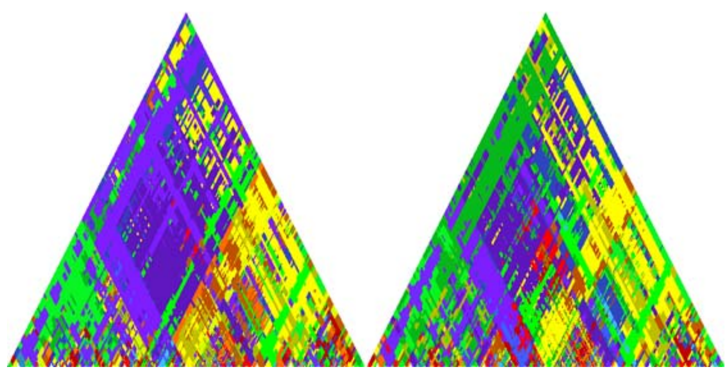

Fig. 28. Keyscape plots for Anton Webern's Variations for Piano, Op. 27, mvmt. 1. Aarden weights used on the left and Krumhansl weights on the right.

5.4 Analysis of Anton Webern's Piano Variations, Op. 27, first movement

The first movement of Anton Webern's Piano Variations, Op. 27 (1936) is given as a final example of keyscape pictures. Anton Webern (1883-1945) was an Austrian composer and a student of Arnold Schoenberg, the developer (around 1920) of twelve-tone music. In twelve-tone music, all twelve chromatic pitches of the octave are played in sequence before any notes are repeated. The goal is to destroy all perception of key and tonality: since no particular note is favored in the music, there can be no tonal center in the music.

Figure 28 shows keyscape plots of the first movement of the Piano Variations. In these plots, it is easy to see the esthetic of destroying a tonal center. As the analysis window becomes larger from bottom to top in the plot, the color regions remain fragmentary. No overall key region becomes dominant in the large-scale structure of the music. This is, of course, intentional in twelve-tone music. Previous example keyscape plots in the article were derived from tonal music, and there has always been a simplification in the key structure towards the music's top-level organization.

Not only are the higher-level regions of the keyscape plot fragmentary, but distant key regions are displayed adjacent to each other. The pitch-to-color mapping based on the circle of fifths also helps to demonstrate the lack of key in this example. Distantly related key regions are represented by contrasting colors. For example, violet and yellow are opposite colors representing key regions of A and F respectively; green and red represent musical regions C and E.

BIBLIOGRAPHY

AARDEN, B. 2003. Melodic expectancy. Ph.D. dissertation. The Ohio State University, 2003.

KRUMHANSL, C. 1990. *Cognitive Foundations of Musical Pitch*. Oxford University Press, New York.

LERDAHL, F. AND JACKENDOFF, R. 1983. *A Generative Theory of Tonal Music*. MIT Press, Cambridge, MA.

NARMOUR, E. 1977. *Beyond Schenkerism: The Need for Alternatives in Music Analysis*. Chicago University Press, Chicago, IL.

SAPP, C. 2001. Harmonic visualizations of tonal music. In *Proceedings of the International Computer Music Conference* (ICMC, Havana, Cuba, 2001), 423-430.

SEGNINI, R. AND SAPP, C. 2005. Scoregram: Displaying gross timbre information from a score. In *Proceedings of the Symposium Series on Computer Music Modeling and Retrieval* (CMMR, Pisa, Italy, Sept. 26-28, 2005).

TEMPERLEY, D. 2002. *The Cognition of Basic Musical Structure.* MIT Press; Cambridge, MA.

Received June 2005; accepted August 2005.