Musical Motion at Different Scales: Creative Analysis and Resynthesis of Musical Contour Spectra

Marc Evanstein
Department of Music, University of California, Santa Barbara, USA
marc [at] marcevanstein.com
http://www.marcevanstein.com

Proceedings of Korean Electro-Acoustic Music Society's Annual Conference (KEAMSAC2020)
Seoul, Korea, 29-31 October 2020

In the context of music technology, Fourier analysis is generally applied directly to sampled sound waves, with the goal of revealing timbral information about the sound or sounds in question. By contrast, this paper presents a software tool ("Spectral Musical Contour Explorer") for applying Fourier analysis to more abstract musical time series; for instance, one can analyze a melody as a time series of pitches, or a recording as a time series of RMS volume measurements. Such analyses can uncover salient and musically meaningful periodicities within the structure of musical works. Moreover, the different time scales of these periodicities reflect the multilevel nature of musical structure (e.g. meter, phrase, form). Finally, the software can be used creatively to resynthesize new pitch and volume contours from a hand-selected portion of the analysed spectrum. In particular, we discuss several compositions by the author that use this process to generate novel musical material from melodic and dynamic contours found in canonical repertoire.

Keywords: Fourier Analysis, Resynthesis, Contour, Composition, Software, Form, Melody.

The mathematical tool of Fourier Analysis is used in a wide range of fields and contexts, and can be applied both to time-series data, or to data distributed over another (e.g. spatial) dimension. In music and audio analysis, the most typical use of Fourier Analysis is the application of a Discrete Fourier Transform (DFT)—most typically a Fast Fourier Transform (FFT)—to a sequence of audio samples. This is a powerful tool for timbral analysis, filtering, synthesis, efficient convolution reverb, and many other musical applications (Smith, 2007).

Aside from forming the basis of many of the software tools that musicians use today, this application of Fourier Analysis played a central role in aesthetic movements within the field of music composition, in particular the advent of so-called *spectralist* approaches in the 1970's (Moscovich, 1997). Within the field of music theory, there has recently been a resurgence of interest in a different usage of Fourier Analysis, namely its application to abstract musical structures, such as pitch-class distributions (Amiot, 2017; Quinn, 2007). Here, rather than samples evenly spaced in time, the analysis focuses on samples within circular pitch-class space, an 'outside-time' musical structure, to use the language of composer Iannis Xenakis (1992).

This paper considers a third category of usage: the application of Fourier Analysis to 'in-time' structures that represent changes in abstract musical parameters, such as pitch and volume. (These fluctuations will be termed 'musical contours' throughout this paper.) Such an approach has been considered sporadically, for example by Nettheim (1992) and Voss (1978), and a similar analysis and resynthesis approach using wavelet analysis has been taken by Kussmaul (1991), but it has yet to receive widespread recognition.

This paper both revives this line of research and presents a newly developed piece of software—entitled "Spectral Musical Contour Explorer"—for creatively analyzing and resynthesizing novel melodies and dynamic contours using this approach. Finally, the creative results of this exploration are discussed, in the form of several of my own compositions.

Non-technical Introduction to Fourier Analysis

For the sake of readers coming from a more musical than mathematical background, we begin with a non-technical introduction to the tool of Fourier Analysis, as applied to a recorded waveform. Fourier Analysis is what we use when we talk about the spectrum of a complex sound; for instance, when we say that a clarinet tone has only odd harmonics, or that the first harmonic of a trumpet is stronger than its fundamental, we are referring to the results of Fourier analysis. The central idea is that a complicated motion—in this case, the motion of an air particle under the influence of a trumpet or clarinet—can be decomposed into a superimposition of very simple motions at different speeds.

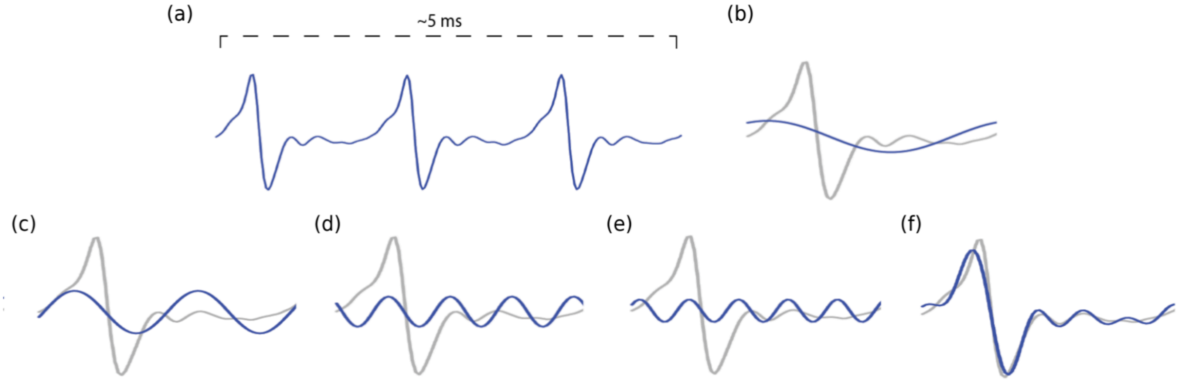

Figure 1. (a) Short excerpt from a trumpet waveform showing a repeated fluctuation in air pressure. (b)-(e) One period of that fluctuation (light gray), with the 1st (fundamental), 2nd, 4th, and 5th harmonics isolated, respectively (blue). (f) The recombination of those harmonics (blue) as compared with the original waveform (light gray).

Figure 1a shows the waveform (i.e. graph of the fluctuation in air pressure) of several periods of a trumpet tone. The unique shape of the waveform creates the trumpet's sonic signature. Note that these fluctuations happen very quickly; the shape repeats three times over the course of 5ms, which translates to 600 oscillations per second (Hz), or roughly a concert D5.

What Fourier synthesis does is break down this complex signature into a sum of sine waves at integer multiples of the 600Hz frequency of the complex pattern. Thus, for a 600Hz trumpet tone, we have components at 600Hz, 1200Hz, 1800Hz, 2400Hz, etc., which we would term the 1st, 2nd, 3rd, and 4th *harmonics* or *partials*. Each partial has its own weighting (*amplitude*) and alignment (*phase*).

We can see what this looks like in Figures 1b-e, which show the 1st, 2nd, 4th, and 5th partials of the trumpet waveform respectively (the original waveform is shown in gray for reference). The first partial completes one cycle for every cycle of the complex trumpet tone; thus it, like the trumpet tone, is oscillating at 600Hz. The second partial completes two cycles for every cycle of the trumpet tone; thus it is oscillating at 1200Hz. Of the four partials shown, notice that the second partial is the strongest, and that the phase of each of the sine waves is such that its peaks and valleys align well with the peaks and valleys of the complex waveform.

Figure 1f shows the sum of these sine waves, which very nearly reproduces the original trumpet waveform. In fact, there is no physical or acoustical difference between the simultaneous sounding of the sine waves in 1b-e and the sound of their sum in 1f. If we wished to reproduce the original trumpet wave with perfect fidelity, we would simply need to include the remaining relatively weak higher harmonics.

It turns out that there is only one way to break a complex wave shape into a sum of sine waves like this, and we call this unique combination of harmonics with different amplitudes and phases a *spectrum*. Thus, when we say that the sound of a trumpet has a strong second harmonic, we mean that the effect of the complex pressure wave produced by a trumpet is identical to the effect of a very specific combination of sine waves added together, and that the second of these sine waves is the strongest.

Musical Countour Spectra

In a musical context, the term *spectrum* is very readily associated with *timbre* and with the direct application of Fourier Analysis to a recorded waveform. Indeed, many canonical examples of the "Spectral" music that emerged in the 1970's (e.g. Gerard Grisey's Les Espaces Acoustiques) are based specifically on transcribing the results of such an analysis to music notation (Féron, 2011).

However, because Fourier analysis is an abstract mathematical tool, it can be just as easily used to analyze the variation in any other musical parameter, at any time scale. For instance, the pitch of a melody can be seen as a time-varying property, operating on the scale of seconds rather than milliseconds. Such analysis, combined with creative resynthesis, can be a source of novel musical material, as we shall see.

Figure 2 depicts the melody of "Pop Goes the Weasel," first in traditional music notation, and then as a time-varying pitch contour. By depicting the pitch of the melody in this way, we see that it is, mathematically, just like the trumpet waveform from before. The only difference is that this wave represents the motion of an abstract musical parameter, rather than of air pressure directly, and that the variation is on the scale of seconds rather than milliseconds.

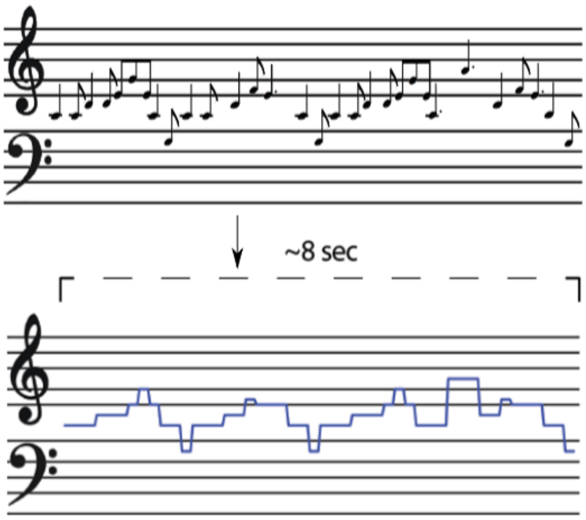

Figure 2. The melody "Pop Goes the Weasel," first in traditional music notation, and then reinterpreted as a time-varying pitch contour.

There is, therefore, no reason we cannot apply Fourier analysis to this pitch contour, just like we did with the trumpet waveform. As with the trumpet, these oscillations operate at 1x, 2x, 3x, etc. the frequency of the melody itself[1], and each of these 'partials' has its own amplitude and phase, with some partials being especially influential.

Figure 3 shows some of the lower partials. The first partial (Figure 3a) is not particularly strong, but its phase is nevertheless aligned so that the peak coincides with the highest note (A4) of the melody. The same can be said of the second partial (Figure 3b). The strongest component is the fourth partial, which completes four full cycles over the course of the melody (Figure 3c). Why is this?

The reason is that the melody itself is in four parts, and each of its first three phrases follows the same pattern of low then high. The final phrase, starting on the A, the highest note in the melody, is somewhat of an exception. In order to compensate, the first and second partials are aligned so as to peak at this exact moment, as is the eighth partial shown in Figure 3d. The eighth partial also helps to create the more local peaks at G4 in the first and third phrases.

Adding together the partials depicted in Figure 3a-d, we arrive at the contour shown in Figure 4a, which tracks the motion of the melody fairly faithfully, albeit a little too smoothly. In order to achieve the flat pitch plateaus that our western ears have come to expect, we need to include more rapid fluctuations like the 20^{th} partial (Figure 3e) to help flatten out the peaks of the slower sine waves (Figure 4b). As with the trumpet waveform, by including enough partials we can reproduce the original melodic contour with perfect fidelity.

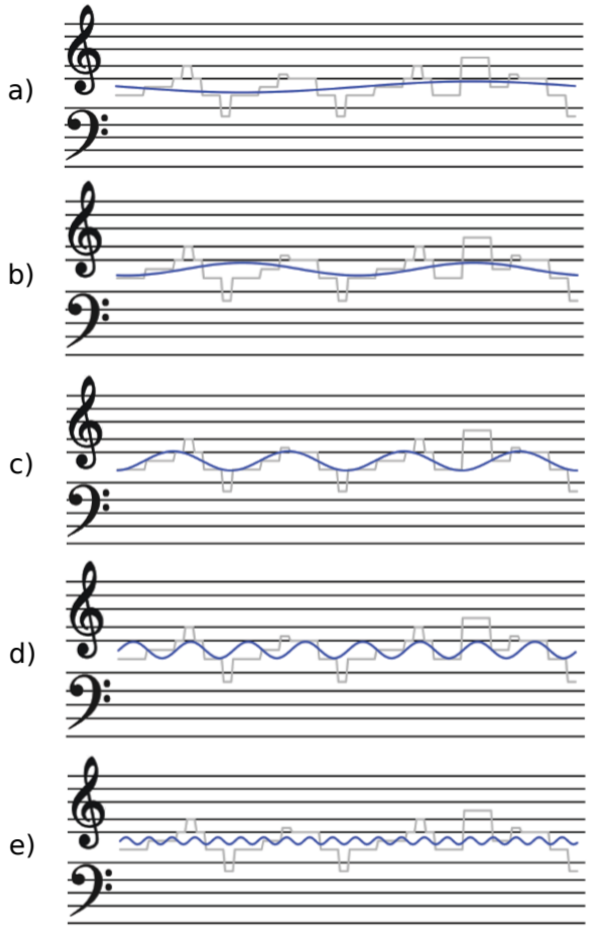

Figure 3. The (a) 1^{st}/Fundamental, (b) 2^{nd}, (c) 4^{th}, (d) 8^{th}, and (e) 20^{th} harmonics of harmonics of the pitch contour of the melody "Pop Goes the Weasel" (blue), superimposed on the original contour (light gray).

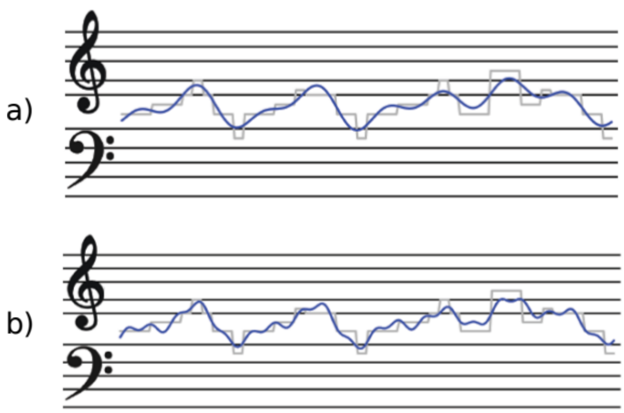

Figure 4. (a) The sum of partials 1, 2, 4, and 8, producing a passable approximation of the melodic contour. (b) Sum of partials 1, 2, 4, 8, and 20, showing that partial 20 helps to flatten out some of the peaks into plateaus.

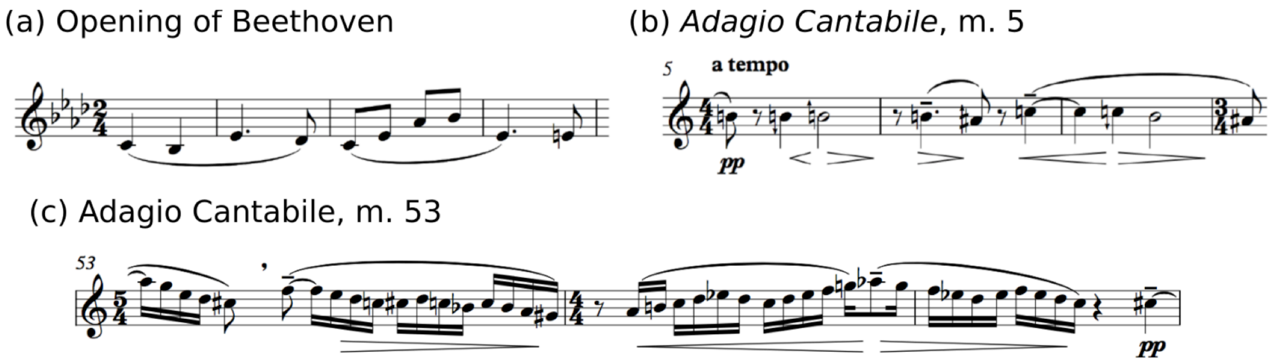

Figure 5. Comparison of (a) the opening melody from the second movement of Beethoven's Pathétique and (b) a corresponding passage in Adagio Cantabile. (c) is an excerpt from towards the end of the work.

A Mathematical Schenkerian Analysis

The above should give some indication of the potential for using this kind of Fourier Analysis as an analytical tool. Among the magnitudes and phases of the various partials was valueable information about the structure of the melody, from its overall shape (1st and 2nd partials), to its phrases (4th partial), to hints of its motivic and rhythmic structure (8th partial). Those familiar with the theories of Heinrich Schenker may note a certain kinship here, in that both Schenkerian analysis and this application of Fourier Analysis represent a hierarchical view of musical structure. (In Schenkerian analysis, this hierarchy is represented by a range of interrelated structural levels, from background (*Ursatz*), to middleground, to foreground (Cadwallader & Gagné, 2007).)

From an analytical point of view, the process described above may provide a valuable complement to the process of Schenkerian Analysis, with the former valued for its objectivity, and the latter for its subjectivity.

A Tool for Creative Resynthesis

The illustrations in Figures 2-4 were produced using a tool that I created for analysis and resynthesis of musical contour spectra, called "Spectral Musical Contour Explorer." This program was created in Python, using PyQt5 as the underlying GUI framework, and using an embedded ChucK (Wang & Cook, 2004) binary for rudimentary sound synthesis. A more complete screenshot of the program is shown in Figure 5.

To begin with, the user is allowed to load either a MIDI file or a WAV file. In the case of a MIDI file, the average pitch of all active MIDI notes over time is plotted against a grand staff[2], while in the case of a WAV file, a plot of the variation in RMS volume over time is displayed. When inputting a MIDI file, the user is prompted for a length, in quarter notes, to assign to each sample; when inputting a WAV file, the user is prompted for the desired window size for calculating the RMS.

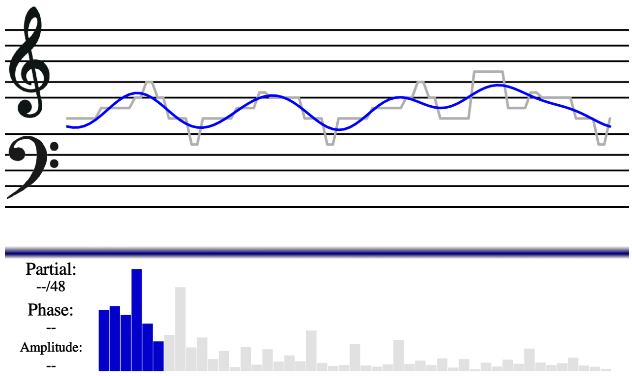

Figure 6. Screenshot from "Spectral Musical Contour Explorer." The bottom half of the screen shows the contour spectrum of the melody, with active partials in blue and inactive partials in gray. The top part of the screen plots both the original melody (in gray) and the reconstructed melody (in blue).

By mousing over the partials of the spectrum in the bottom half of the screen, additional information about phase and amplitude can be viewed, and the user can then click any these partials to toggle them on and on or off. In this way, any partial spectral reconstruction of the contour can be achieved. Finally, the samples of the reconstructed contour can be exported in the form of a text file, so that they can be used in a composition, or for further analysis.

Creative Results

Adagio Cantabile

The first piece in which I made use of this technique was *Adagio Cantabile*, for oboe and guitar. Using as source material the main theme of the second movement of

Beethoven's *Sonata Pathetique, Op. 13*, I performed Fourier analysis and resynthesis on both pitch and rhythm independently (encoding rhythm as a sequence of note length samples[3]). A happy accident occurred in this process: since I was treating pitch values as continuous, rather than discrete, I ended up with microtonal inflections in the resulting resynthesis. This ended up becoming a central aspect of the oboe part.

After exploring the space of possible reconstructions, both in pitch contour and rhythm, I ended up with a collection of short melodic snippets, which I ultimately assembled using pencil and paper. Figure 6 compares the opening melody of the Beethoven (a) with a two excerpts from *Adagio Cantabile* featuring a partial reconstruction of the melody. In (b), the overall contour of the melody has been removed, leaving only the slight microtonal deviations. In (c), taken from near the end of the work, it is the local ornamentation—the higher frequency information—that has been removed, leaving a melodic line that sweeps gradually up and down. In this latter case, I allowed myself considerable flexibility in choice of accidentals, letting my ear guide such decisions intuitively.

Unraveled

The second piece in which I used this technique was *Unraveled,* for Percussion Quartet and Impossible Electronic Orchestra. The title gives a hint as to the source material: the famous melody from Ravel's Bolero. I used the software described here to analyse and resynthesize the melody in various degrees of recognisability, and then had these reconstructions performed by an "Impossible Orchestra," consisting of pitch-bent samples of orchestral instruments.

As with *Adagio Cantabile*, then, the contour in question is a melodic pitch contour. An added wrinkle in this case, however, is that rolls in percussion parts are used to emphasize the individual partials of the melodic contour, with many of these rolls superimposed on one another at a given time. Thus, though the fission process of Fourier analysis, the monophonic melody gives rise to a heterophonic accompaniment, one which emphasizes details within the melody itself.

Anamnesis

The third (and most recent) work that I composed using spectral analysis and resynthesis of musical contours is *Anamnesis* for Chamber Orchestra. Anamnesis differs from *Adagio Cantabile* and *Unraveled* in that the musical contour being analysed was a dynamic contour, rather than a pitch contour. Here, again, I used a famous work as the source material for analysis: the *Allegretto* from Beethoven's *Symphony No. 7*.

Figure 7a shows a screenshot (from the program *Audacity*) of the recording by Carlos Kleiber and the Vienna Philharmonic Orchestra. Figure 7b shows this same recording, as loaded into "Spectral Musical Contour Explorer": a dynamic contour has been created by calculating RMS values for every half-second of audio, effectively resulting in a unipolar version of what we see in Audacity. Figures 7c, d, and e show three examples of strong periodicities found in the dynamic contour through analysis. Notice, for instance, how the second harmonic depicted in Figure 7c highlights the two main peaks of intensity within the movement.

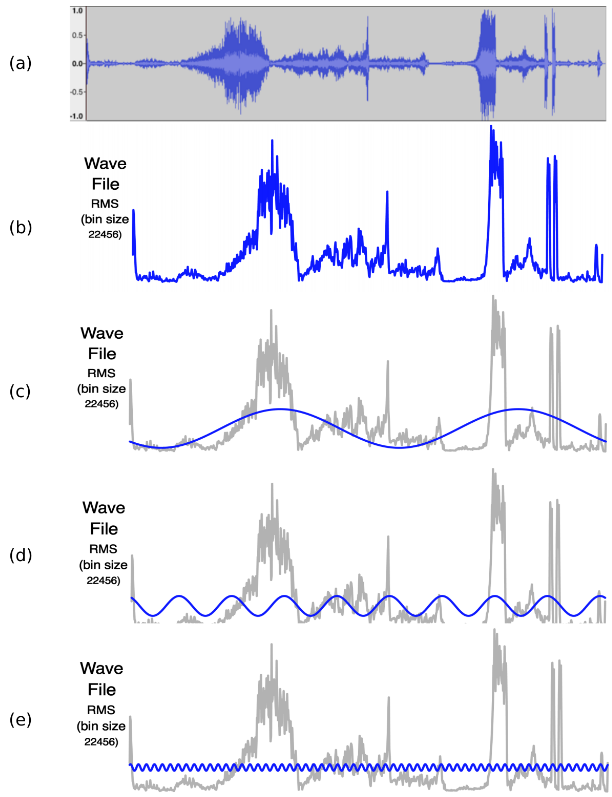

Figure 7. Analysis of the dynamic contour of the *Allegretto* from Beethoven's *Symphony No. 7*. (a) The dynamic shape of the movement as shown in Audacity. (b) The same contour, as represented in "Spectral Musical Contour Explorer." (c), (d), and (e) Three prominent periodicities found at different time scales.

I then used the *SCAMP* libraries for computer-assisted composition in Python (Evanstein, 2018) to orchestrate these different layers of motion, with some instruments playing the larger swells, others playing the mid-level swells, and still others playing the fastest-moving swells. Thus, as in *Unraveled*, Fourier Analysis broke a single contour (this time a dynamic contour) into a heterophonic texture of simple gestures. Figure 8 shows an excerpt of the texture in the violins from the opening of the

Figure 8. String excerpt from the opening of *Anamnesis*.

work, consisting of many short, overlapping swells. Below, one can see a larger swell beginning, tremolando, in the cello section.

It should be noted that, as in Adagio Cantabile and Unraveled, the process of musical contour analysis and resynthesis was merely the starting point for the composition. The final work also resulted from numerous other musical processes and decisions, which were largely intuitive in nature.

Conclusions

There are several possible avenues of further research and compositional practice. From the point of view of composition, each new contour represents a different initial condition for the creative process, as does each possible approach to resynthesis (e.g. removing all but the low partials, the odd partials, the prominent partials, etc.). As the above examples illustrate, this approach can generate snatches of musical material (as in *Adagio Cantabile*) and/or it can form the basis of the work's overall form (as in *Anamnesis*).

Possibilities can be further expanded by the development of the tool itself. For example, one could incorporate contours based on other musical parameters: instead of RMS volume, an inputted sound file could be analyzed in terms of its variation in spectral centroid, spread, flux or kurtosis, or on its zero-crossing rate.

Another interesting possibility would be to allow for modification of the phase of partials before resynthesis of the contour. In many cases, the effect of phase is as important as, if not more important than, magnitude in establishing structural boundaries within a data set (Bartolini et al., 2005).

In short, I envision this approach, and the tool I have developed, as one among many that could serve as a source of inspiration for composers in their creative process.

References

Amiot, E. (2017). A Survey of Applications of the Discrete Fourier Transform in Music Theory. In G. Pareyon, S. Pina-Romero, O. A. Agustín-Aquino, & E. Lluis-Puebla (Eds.), The Musical-Mathematical Mind: Patterns and Transformations (pp. 17–28). Springer International Publishing. https://doi.org/ 10.1007/978-3-319-47337-6_3

Bartolini, I., Ciaccia, P., & Patella, M. (2005). WARP: Accurate retrieval of shapes using phase of Fourier descriptors and time warping distance. IEEE Transactions on Pattern Analysis and Machine Intelligence, 27(1), 142–147. https://doi.org/10.1109/TPAMI.2005 .21

Cadwallader, A. C., & Gagné, D. (2007). Analysis of Tonal Music: A Schenkerian Approach. Oxford University Press.

Evanstein, M. (2018). SCAMP: a Suite for Computer-Assisted Music in Python. GitHub. https://github.com/ MarcTheSpark/scamp

Féron, F.-X. (2011). The Emergence of Spectra in Gérard Grisey's Compositional Process: From Dérives (1973–74) to Les espaces acoustiques (1974–85). Contemporary Music Review, 30(5), 343–375. https://doi.org/ 10.1080/07494467.2011.665582

Kussmaul, C. (1991). Applications of the Wavelet Transform at the Level of Pitch Contour. PROCEEDINGS OF THE INTERNATIONAL COMPUTER MUSIC CONFERENCE, 483–483.

Moscovich, V. (1997). French Spectral Music: An Introduction. Tempo, 200, 21–28. https://doi.org/10.1017/S0040298200048403

Nettheim, N. (1992). On the spectral analysis of melody. Interface, 21(2), 135–148. https://doi.org/10.1080/09298219208570604

Quinn, I. (2007). General Equal-Tempered Harmony: Parts 2 and 3. Perspectives of New Music, 45(1), 4–63. JSTOR.

Voss, R. F. (1978). ''1/f noise'' in music: Music from 1/f noise. The Journal of the Acoustical Society of America, 63(1), 258. https://doi.org/10.1121/1.381721

Smith, J. O. (2007). Mathematics of the Discrete Fourier Transform (DFT): With Audio Applications. Julius Smith.

Wang, G., & Cook, P. (2004). ChucK: A programming language for on-the-fly, real-time audio synthesis and multimedia. Proceedings of the 12th Annual ACM International Conference on Multimedia, 812–815. https://doi.org/10.1145/1027527. 1027716

Xenakis, I. (1992). Formalized Music: Thought and Mathematics in Composition. Pendragon Press.

[1] For those more familiar with Fourier analysis, it will be apparent that I am using a window size equal to the whole length of the melody. There is, therefore, an underlying assumption that the melody itself is cyclic. This may be more or less appropriate in different situations.

[2] In the Pop Goes The Weasel examples above, the inputted MIDI file was monophonic.

[3] The results turned out to be quite interesting with rhythm: When no frequencies (except DC) of oscillation were present, the rhythm was static, with all notes the same length. When lower partials were included, the rhythm started to accelerate and decelerate at the faster and slower parts of the melody. As I included faster and faster oscillations, these accelerandi and decelerandi became more and more local, until all of the detail of the original rhythm was recreated.