Spectral Melody Decomposition

"Spectral Melody Decomposition" is simply the term I use for applying Fourier analysis to a melodic contour, as opposed to its more typical musical use in analyzing the frequency content of sound waves directly. By applying Fourier analysis to this more abstract musical information, we can uncover interesting information about what makes a melody tick, and how it evolves at different time scales. In my own compositional practice, I find it fruitful to decompose the contours of great melodies and then reassemble them in creative ways, yielding novel results.

If that makes no sense, read on. The point of this paper is to make sense out of that last paragraph.

FOURIER ANALYSIS

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 the fundamental, we are referring to the results of a Fourier analysis. The basic 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 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, or roughly a concert D5.

What Fourier synthesis does is break down this complex signature into a sum of simple motions – sine waves, in fact – 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 grey 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 each of the sine waves is positioned so that its peaks and valleys align as well as possible with the peaks and valleys of the complex waveform.

Figure 1: Trumpet Note

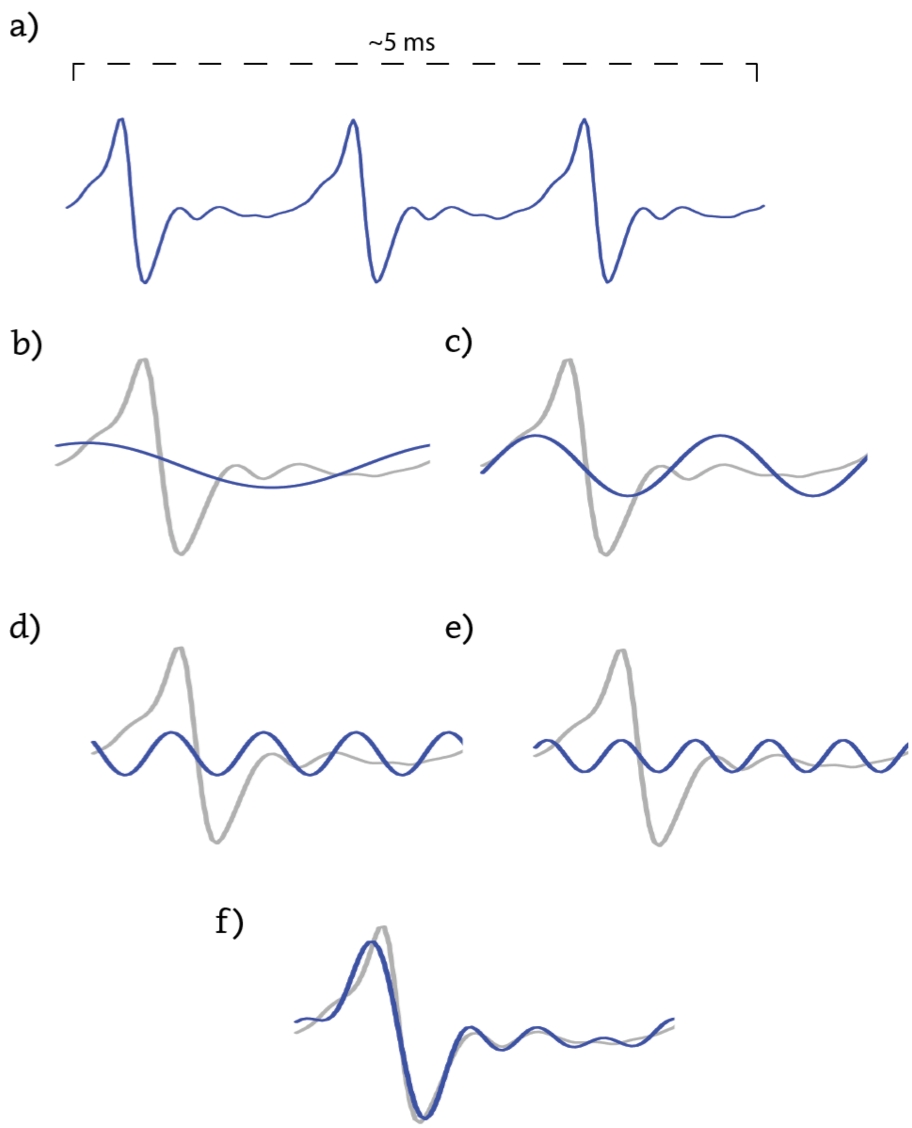

Figure 1f shows the sum of these sine waves, which (as if by magic) very nearly reproduces the original trumpet waveform. One of the remarkable properties of waves is that they add together in the medium in which they propagate. For instance, if an air particle is under the influence of two waves, one of which is pushing it to the right with a force of 5 and the other to the left with a force of 3, it responds by accelerating to the right under a net force of 2. For this reason, there is no physical or acoustical difference between the simultaneous sounding of the simple sine waves in 1b-e and the sound of their sum in 1f. If we want to take it further and reproduce the original trumpet wave with perfect fidelity, we simply need to include the wide variety of relatively weak higher harmonics.

It turns out that there is only one way that we can add together sine waves in this way to produce any given complex wave shape; the spectrum is unique. Thus, this is what we mean when we say that the sound of a trumpet has a strong second harmonic: the effect of the complex pressure wave produced by a trumpet is identical to the effect of a multitude of carefully tuned sine waves added together, and this unique spectrum has a strong second harmonic. Sonically, this translates to a certain type of brightness in the sound.

FOURIER ANALYSIS OF MELODIC CONTOUR

What's wonderful about Fourier analysis is that it is a mathematical abstraction, so while in this case it tells us about the spectral components of a fast oscillation in air pressure (i.e. a sound wave), it can be just as easily used to analyze the variation in any other parameter at any time scale.

Figure 2a shows the melody for “Pop Goes the Weasel” in traditional music notation. As you can see, the pitch varies up and down over time. By replacing the noteheads with a line, we can see in figure 2b that the pitch contour is a wave like any other wave. The only difference from the trumpet wave above is that this is a wave representing the motion of an abstract parameter (i.e. pitch, rather than the air pressure directly), and that the variation is on the scale of seconds rather than milliseconds.

Since the pitch contour is a wave, there is no reason that we cannot apply Fourier analysis to it, 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.

Figures 2c-f show some of the lower partials. The 1x frequency, shown in figure 2c, is not particularly strong, but you can see that the phase is aligned so as to make the peak coincide with the highest note (A4) of the melody. The same can be said of the 2x frequency (figure 2d).

The strongest component is the one that makes 4 cycles over the course of the melody (figure 2e), and when you think about it, this is no surprise at all, as it expresses the 4-part structure of the melody! Each 2-bar phrase (in 6/8) starts low in the first bar and peaks at the beginning of the second bar.

There is one exception, of course: the final phrase, whose downbeat is the highest note of the melody. The 1x and 2x components have provided a peak at this spot, but we can also see from figure 2f that the 8x component provides a boost here. The 8x component also helps to create the more local peaks at G4 in the first and third phrases.

[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.

Figure 2: Pop Goes the Weasel

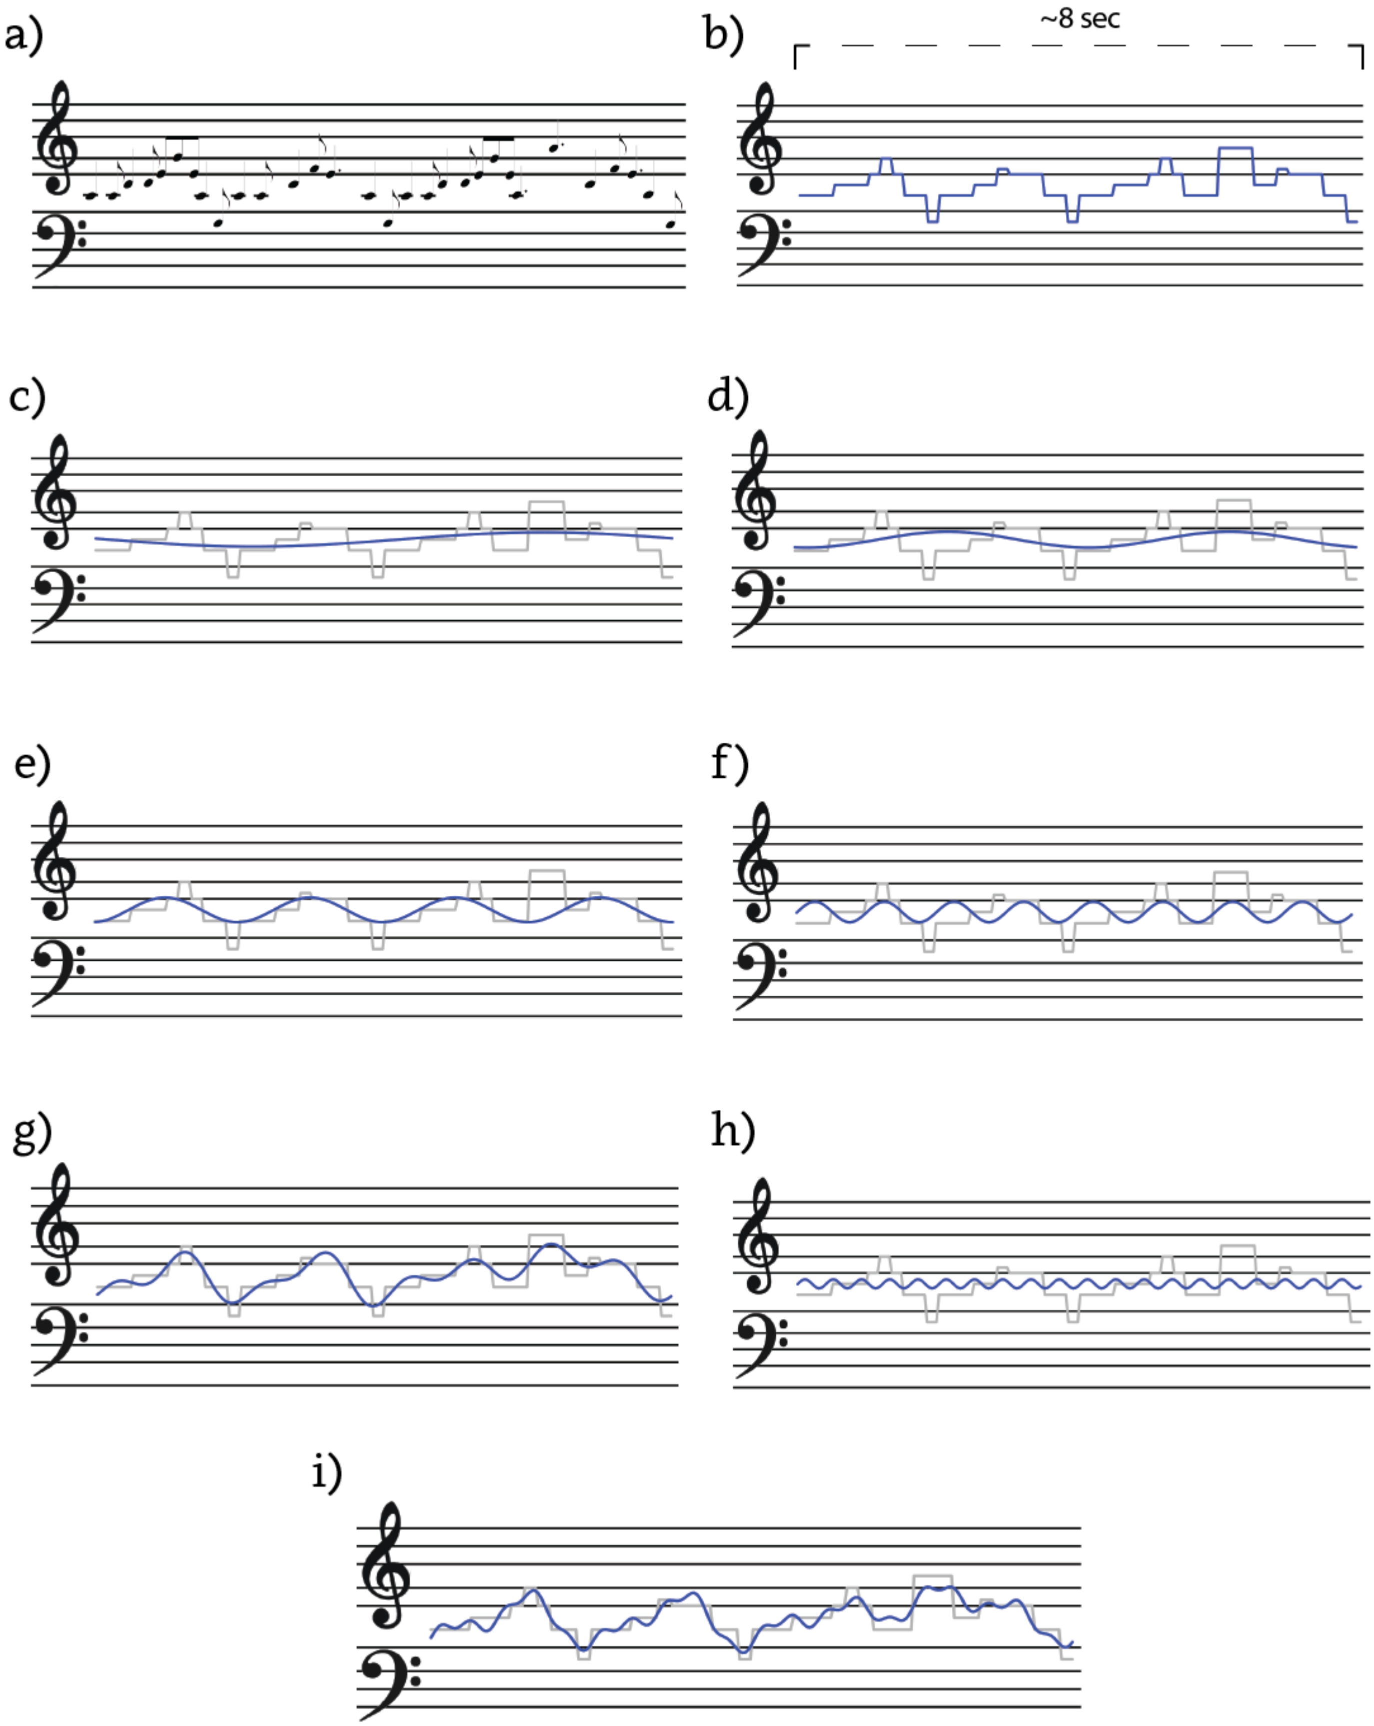

Adding together several of these slower moving components, we arrive at the contour in figure 2g, which tracks the motion of the melody fairly faithfully, albeit a little too smoothly. (After all, pop goes the weasel is not performed with continuous glissandi.) In order to achieve the flat pitch plateaus that our western ears have come to expect, we need higher frequency components like the 20x oscillation in Figure 2h. Figure 2i adds this in, and you can see that it is carefully calibrated to counteract some of the bulges of the slower moving sine waves. Once again, by including enough partials, we can reproduce the original melody contour with perfect fidelity.

To sum up, we have seen that the application of Fourier analysis to the pitch contour a melody reveals a set of simple undulations that combine to form a complex impression of tension and release.

SCHENKERIAN THOUGHTS

Those familiar with the theories of Heinrich Schenker may note a certain kinship between Schenkerian analysis and the ideas presented above. After all, both reflect the multiscale nature of musical structure. When we use Fourier analysis to investigate the pitch contour of a melody, the lower partials that result are responsible for the overall shape of melody, while the higher partials are responsible the more surface-level ornamentation. These are analogous to the background and foreground levels of a Schenkerian analysis, respectively.

COMPOSITION PRACTICE

I have been investigating various ways of using the spectral decomposition of melodic/musical contours in my compositional practice. Among the possibilities I have been exploring are:

1) Experimenting with partial reconstruction of a melody. What does it sound like to take a melody and remove the slower moving components? What about the faster moving ones? I have used this approach in my piece, *Adagio Cantabile*, which I discuss below.
2) Rephasing: What does it sound like if we simply change the phase of various components without changing their amplitude? Is the new melodic contour interesting?
3) Accompanying a melody with swells corresponding to the most important components of its pitch contour.

4) Turning a melodic contour into a polyphonic texture by partitioning its components into groups and having different lines follow the contours of different groups.
5) Applying Fourier analysis to parameters other than pitch. For example, by applying Fourier analysis to the dynamic envelope of a performance, we get swells at different frequencies that can combine to recreate the original dynamic contour. (The beauty of using volume as the parameter of analysis is that amplitudes add.)
6) Employing Fourier analysis on one parameter, and then using the resulting spectrum to shape musical material in an entirely different parameter. To what extent does a gesture remain convincing / meaningful when translated to a different parameter?

What makes all this so interesting to me is that there is a kind of alchemy involved in the creation of a good melody. Why is it that in some contexts a climactic high note fills us with exultation, whereas in other contexts it leaves us cold? I think that a partial answer has to do with the richness of the "melodic contour spectrum": when several interlocking cycles, peaking at different times, finally rise up together at the same moment, there is something in us that responds.

Of course that hardly gives a formula for a good melody; the only real "formula" for writing a good melody is to build up one's intuition over years of study, immersion and musical practice. The interest for me as a composer is rather to take the contour spectrum of a famous and/or personally meaningful melody and apply it to the shaping of new material for a new piece. The goal is to discover a world of possibilities buried within the original melody.

Every piece I write using this technique is unique, and so demands a different way of using the tool. Also, depending on the nature of the piece, I may use the results more strictly, or more freely. Inevitably, though, the piece will have a mind of its own; it's a matter of finding the most effective collaboration between process and intuition. Mathematical processes can open up amazing and unexpected worlds, but when mindlessly employed, they can also lead to exceedingly boring and ineffective music.

ADAGIO CANTABILE

Adagio Cantabile is an example of the first approach in the list above, namely experimenting with partial reconstructions of a melody. Using the melody from the second movement of a Beethoven sonata as my model, I tried removing different components of the pitch contour to see what would result.

When the lower "partials", i.e. the larger, slow-moving components, are removed, the melody ceases to have overall highs and lows, instead wiggling subtly around a central pitch (hence, the microtonal inflection in the oboe part). The most extreme example from the score is found in measures 5-12; note the tight motion around B4 (Figure 3b).

Figure 3

a.) Original melody

b.) *Adagio Cantabile*, m. 5

c.) *Adagio Cantabile*, m. 48

Gradually, as the piece progresses, more and more of these slower moving partials are included, and the oboe starts to skirt around the edges of the original Beethoven, with reminiscent gestures here and there. The point at which the oboe comes closest to the original melody occurs in m. 48, though even then the reference it's unlikely to be recognized (Figure 3c). This is intentional; the point of this process is to open up a parallel universe, one that owes its existence to Beethoven, but which is only weakly referential.

Toward the end (starting in measure 52), I explore the other side of the contour spectrum, with undulating scales that follow the contour of the melody with the faster-moving components removed.

The guitar part stems from a similar process as the oboe part, except that the process has been applied to multiple lines simultaneously. The original Beethoven is in 4-part harmony, with the oboe based on the top part and the guitar based on the bottom three parts. The guitar proceeds through a series of interludes starting in m. 13, m. 23, m. 33, and m. 40, with the first being based on just the slowest contour components of the accompanying lines, while each subsequent interlude adds faster components. Notice how the first interlude simply consists of the three lines rising up and down once, where as the later interludes have more nuanced contours.

The harmony that results is still largely triadic, since the spacing of the parts stays fairly close to the original; however, the lack of the higher contour components tends to lead to chromatic interpolations. (For instance, a sudden leap of a fourth in the original requires a high-frequency contour component that, if removed, results in a slower chromatic motion up a fourth.)

I should note that, despite the precise scientific manner of the preceding discussion, the actual compositional process was far from clean and orderly. For instance, although the oboe material came from the process I have described, I reworked and remixed it heavily based on my musical intuition. Likewise, while the contours of the guitar interludes were determined mathematically, the rhythmic patterning and the exact pitches of the compound line were chosen intuitively.

A PERSONAL NOTE

It is such a funny thing to write about my process in this way, because it comes off so much dryer and more calculating than it is felt or intended. The truth is, there's a great deal of sentimentality at the root of all this.

In *Adagio Cantabile*, for instance, I chose the Beethoven because it was one of my grandmother's favorite pieces. She used to play it on the piano, and during the last several years that she was alive, I would play it for her every time I visited. She had some form of dementia, and it was amazing to see how hearing the piece would bring her to life: she actually closed her eyes and moved her fingers in mid air as she listened.

This is more than an intellectual exercise; it is an investigation into the music that moves me, with the ultimate goal being to create more music with a similar power to move.