

# Sonifying Letter Frequencies: Generating Chord Sequences from Text

Donald Spector

Dept. of Physics, Hobart and William Smith Colleges, Geneva, NY, USA; spector@hws.edu

## Abstract

As part of an ongoing project to consider the representation of various structures in music, I explore aspects of turning ordinary text into musical forms. Ultimately, the approach of greatest interest is one inspired, in broad terms, by Huffman coding: creating a sonic encoding that reflects the frequency of the characters. However, rather than associating letter frequency with duration of the corresponding notes (as a Huffman coding might do), we use a method motivated by the *tonalité moderne*: we pick a key, and encode frequent letters with chords that are harmonically closely related to the tonic chord of the key, while less frequent letters are associated with harmonically more distant chords. This paper is intended both as a consideration of how to conceive of an effective sonification of text, and as a proof of concept, presenting the methodology and sample outcomes, so that this methodology might find broader use.

## Introduction

This paper is part of an ongoing project to explore sonification—the representation of data via sounds [3]—for various forms of data, seeking rules in which the sounds generated in some fashion reflect the underlying structure of the data. Interest in sonification ranges from the use of sounds to convey information about a set of data (as in the sound renderings of black hole mergers [5]) to the use of data as a tool to generate aesthetically appealing sounds. (Of course, sonifications can sometimes address both these criteria.) In my recent work in this arena, for example, I developed methodologies to represent the sequence of steps to solve a mathematical puzzle with chords so that the harmonic sequences reach resolution when the puzzle is solved [8], while in a separate paper, I developed methods for presenting the game play of competitive games either melodically or rhythmically in a way that captures some of the ebb and flow of game play, but in which the main goal is to create aesthetically appealing music [7].

Here, I seek to explore whether there are interesting ways to generate musical passages from ordinary English text. Given the large space of data one has to consider—whether the thousands of words or dozens of letters—it seems that such an effort should focus on using text to create musically interesting sounds, rather than trying to provide a useful auditory representation of that text. (In fact, of course, auditory representations of text already exist; one can simply read them aloud, or spell them aloud, but neither of these really captures the spirit of sonification, which should provide qualitatively new insights into or aesthetic experiences from some data.) Thus my goal here is to examine how one might use textual material as the source material to generate (some aspect of) a musical composition. In this paper, my focus is not simply on presenting a particular sonification of text, but also on discussing how to approach the question of determining what sorts of sonification would be compelling. This discussion is informed both by my previous studies and by experimentation in this specific context. Still, I should be clear that the discussion below is not intended to be exhaustive in the insights it offers regarding the pros, cons, and other features of various potential sonification methodologies.

In seeking to sonify a block of text, one needs to consider what data to use as the source, and what characteristics to use to generate sounds. As for the data, the two primary units of data one might employ as the inputs to the sonification rules are the words or the characters. There are two arguments against using

words. The primary problem is that the set of words is too large: there are thousands of relatively common words, so it would be a problem to come up with a sonification rule that meaningfully represents the various elements of this set, especially since there is not a clear organizational property to use as the criterion to generate sounds. The set of words does not have enough structure. Furthermore, if one were to seek a sonification based on semantics, one might argue that setting texts to music when writing songs or operas or the like already does this: the composers are choosing music that is, at least in the composer's view, semantically appropriate. It is true that such mappings do not have the automaticity of a typical sonification rule, but at the very least, it would seem one could argue that a semantically-based sonification would be an extension of existing practices of setting words to music.<sup>1</sup>

Consequently, we turn our focus to developing sonification rules that turn characters into sounds. In general, with sonification, one seeks to identify properties that can be turned into sounds that represent some aspect of the structure of the data set; if we are seeking to encode things using the familiar forms of Western music—the *tonalité moderne*—it is also important not to have too many distinct entities to encode. Thus, in focusing on characters, we choose to limit our exploration to textual passages from which all non-alphabetic characters have been stripped, and with no distinction made between upper and lower case characters, so there are only 26 distinct entities to be sonified<sup>2</sup>

To get a feel for the problem, it is useful to see what various sonification methodologies produce. I here present some of my findings. Attaching different notes to each letter, without any guiding principle, produces music of no apparent structure or aesthetic appeal. If all the notes in the chromatic scale are used, the result is rather formless atonal music that, even for someone like the author who appreciates atonal music, has little to recommend it. An alternative method, then, might be to seek a mapping of notes to letters in a way that all possible notes fit together in some way. Two natural choices are to divide the letters of the alphabet among the five notes of one of the standard pentatonic scales (e.g., the notes  $C\sharp$ ,  $D\sharp$ ,  $F\sharp$ ,  $G\sharp$ , and  $A\sharp$ ) or to divide the letters of the alphabet among the notes of a blues scale (e.g.,  $D$ ,  $F$ ,  $G$ ,  $A\flat$ ,  $A$ , and  $C$ ).<sup>[6]</sup> An advantage of these scales is that they lack some of the dissonances that arise in ordinary major or minor scales. For example, one can safely use any of the notes of the blues scale when any of the basic blues chords (I, IV, or V) are providing the harmony. Nonetheless, using these frameworks proves unsatisfactory. The resulting melodies might be technically acceptable harmonically, but nonetheless feel aimless and formless, even when rhythmic variety is introduced as part of the sonification rules; one finds neither the satisfaction of the music having direction, nor the excitement of capturing a large space of sounds.

The apparent need to find some structure suggests that a sonification scheme based on chords and harmonies is more likely to be effective, so that the music will have some anchors to provide structure. Here, I take some inspiration from Huffman coding [1]. Huffman codes encode characters differently based on how often they appear. The goal in a Huffman code is to maximize the amount of information that can be transmitted via a given number of bits, and so such codes attach short codewords to more frequently occurring symbols and longer codewords to less frequently occurring symbols.<sup>3</sup>

My focus, however, is on tonal sonifications, not rhythmic ones (in some sense, Morse code already provides a rhythmic sonification), which means not taking the Huffman code approach literally, but rather

---

<sup>1</sup>It is possible that there would be something novel to do using a word2vec or similar structure [2], and focusing not on mapping individual words, but on the relative closeness of textually adjacent words. This would be like focusing on bonds or plaquettes in an Ising-like spin system [4] or sonifying draughts by considering the displacement vector of a move rather than the absolute position of a move [7].

<sup>2</sup>It is worth noting that there are hybrid possibilities, such as encoding words based on the number of characters in a word; it is not clear if one should think of this as character-based or word-based. However, based on some limited trials, there did not appear to be enough structure in the sequences of word lengths for a sonification based on this data to yield appealing results.

<sup>3</sup>Huffman codes are not the only such codes that do this; they are just the optimal such codes meeting certain criteria. Shannon-Fano codes also associate greater frequency with shorter codeword length [1], and, in fact, Morse code was also designed so that more frequently used letters would tend to have shorter Morse code representations.

as a way to focus on meaningful structure within a block of text. With that in mind, we are led to explore sonifications in which each letter is associated with a particular chord, but the chord associations are determined by the frequency of those letters. One can do this as follows. First, select a key, and then assign the tonic major chord of that key to the most common letter. Then choose the next most frequent letters to be assigned chords harmonically close to the tonic chords (e.g., the dominant seventh, the subdominant, and the relative minor for the second through fourth most common letters), and then, gradually, as one moves to less and less common letters, one associates chords are harmonically more distant from the tonic chord.

Suppose, then, we have a block of text, and assign chords to each note following a rule as described above. The idea motivating such an approach is that more common letters will, because they occur more often, regularly bring the chord progression back to our harmonic home, but the less frequently occurring letters will provide the harmonic interest of moving away from the most basic chords associated with a given key. Thus the structure within a text of having characters appearing with different frequency becomes a musical structure that is familiar from within western music: harmonic progressions (i.e., sequences of chords) that are strongly centered on the tonic, but explore the possibilities offered by the chromatic scale to reach out to more distant chords. Thus, the variation of letters in text is being used to produce chord progressions whose structure maintains an aesthetically pleasing relationship harmonically.

Happily, this method, which seems reasonable in theory, meets expectations, and provides some interesting chord progressions. Before presenting a concrete example, I want to be clear about the aims of such an exercise. For simplicity, I left rhythm out of this. My view is that this methodology will generate chord sequences, but then composers would be charged with taking those chord sequences and composing on top of them, using their own judgment to determine the number of beats spent on each chord, for example. To be sure, one could construct a rule that determined this for each chord, but there did not appear to be a compelling aspect within the structure of the sequences of characters to come up with a rule that would not feel arbitrary. I also imagine leaving it to a composer to develop suitable melodies to accompany the chord sequences obtained. Thus, one should view this process as a way to use text (and the frequencies of letters within that text) to generate chord sequences that a (human) composer can then use as the foundation for composing music.

Here, I present the methodology by way of an example. To simplify matters, we use an assignment of letters to chords using the standard tabulated frequency tables for letters in English text (rather than calculating the sample for any given text to be sonified). In the example below, we see a reasonable assignment one might use, with the characters listed in order from most to least common within English text, and with the associated chord presented directly below each letter..

E	T	A	O	I	N	S	H	R	D
C maj	G7	F maj	a min	e min	D7	c min	B $\flat$ maj	E $\flat$ dim	g min
L	C	U	M	W	F	G	Y	P	B
d min	B dim	A7	Emaj7	f min	Ab7	a7	Cmaj7	Gmaj7	gm7
V	K	J	X	Q	Z				
F $\sharp$ dim	Fmaj7	bbm7	F $\sharp$ dim	E dim	A maj				

In this particular example, the tonic chord is C major. While in general, you will see that chords earlier in the list are harmonically closer to the tonic, there are some spots farther down the list, such as around the letters Y and P, where there are some chords assigned that are nonetheless somewhat close to the tonic major chord. This was done to create a balance: if we keep all the various ways a note can be associated with a chord (e.g., G major, g minor, G7, g7, and Gmaj7), the tonic, dominant, subdominant, and relative minor will overwhelm the resultant chord sequences available to be assigned to letters, leading to uninteresting results.

On the other hand, it is nice to be able to have, say, G major and g minor triads available. Consequently, it turns out to be an effective aesthetic decision to take some of the chords that are harmonically close to the tonic major, and shift them to relatively uncommon letters. This allows for the chord assignments attached to a block of text to possess the harmonic variety needed to create harmonic interest and aesthetic appeal. And when, on those rare occasions, one of those harmonically close chords assigned to a less frequent letter arises, it helps contribute to the harmonic center without causing the harmonic center to dominate the sonification.

Applying this chord assignment to different texts, one will of course generate different chord sequences. Similarly, one can adopt different choices for which harmonically close and harmonically distant chords are attached to which letters, or what breadth of harmonic distance one wishes to include. (For example in the chord assignment presented above, there are no chords based on C♯.) As an example of the kinds of sequences generated using the above letter-to-chord assignments, the sentence "What is this about?" (the first sentence I tried encoding, for no reason other than a little self-referentiality) yields the chord sequence (using upper case for major and lower case for minor):

f , Bb , F, G7, e, c, G7, Bb, e, c, F, g7, a, A7, G7

In addition to generating chord sequences from different texts, I have also implemented the capability to get text from the web and sonify that, such as from a Wikipedia page. One could imagine, then, the possibility not only of dynamically generated chord sequences, but of finding ways to use such chord sequences to reflect the changing history of an underlying web page. I was not able to find a particularly compelling application of this idea, but it is intriguing enough to warrant mention. Although I worked in English, the same idea can of course be used in other alphabetic languages.

Finally, there is the question of how to think about these chord sequences. As stated above, my conclusion is that the best use for such chord sequences is to treat these outputs as harmonic progressions that a composer might then use as a foundation, using these chord sequences with melodic accompaniments and rhythmic patterns and voicing of the chords within the purview of a composer. This outcome is a bit more limited than the outcomes of some other sonifications, since the musical output is not fully determined, but ultimately it fits within a very familiar model for intersection of mathematics and art: we use mathematical properties to obtain the basic structure around which an artistic creation will be generated, but the precise creation is obtained by a person making deliberately aesthetic choices in the context of that mathematically-derived framework.

## References

- [1] T.M. Cover and J.A. Thomas, *Elements of Information Theory*, 2nd ed., Wiley, 2005.
- [2] For examples, see the Google Code Archive <https://code.google.com/archive/p/word2vec/>
- [3] G. Kramer, ed. *Auditory display : sonification, audification, and auditory interfaces*. Proceedings of the Santa Fe Institute, vol. XVIII. Addison-Wesley, 1994.
- [4] Kramers, H. A., and G. H. Wannier. 1941. "Statistics of the two-dimensional ferromagnet." *Physical Review* 60: 252-262.
- [5] LIGO Scientific Collaboration, <https://www.ligo.caltech.edu/video/ligo20160211v2>
- [6] D.M. Randel, ed. *The Harvard Dictionary of Music*, 4th edition. The Belknap Press of Harvard University Press, 2003.
- [7] D. Spector, "Sonifying Games." In *Proceedings of Bridges 2022: Mathematics, Art, Music, Architecture, Culture*, edited by D. Swart, F. Ferris, E. Torrence, Tessellations Publishing, 2021, pp. 293-296.
- [8] D. Spector, "The Tower of Har(mo)noi." In *Proceedings of Bridges 2021: Mathematics, Art, Music, Architecture, Culture*, edited by D. Riemann, D. Norton, and E. Torrence, Tessellations Publishing, 2021, pp. 257-260.