Progress Towards Recognizing and Classifying Beautiful Music with Computers

MIDI–Encoded Music and the Zipf–Mandelbrot Law

Bill Manaris, Tarsem Purewal, and Charles McCormick
Computer Science Department
College of Charleston
66 George Street, Charleston, SC 29424
(843) 953–8159
{*manaris | purewalt | mccormickc} @cs.cofc.edu*

Keywords: Artificial Intelligence, Computer–Assisted Music Analysis, Zipf–Mandelbrot Law, MIDI.

Abstract

We discuss the application of the Zipf–Mandelbrot law on musical pieces encoded in MIDI. Our hypothesis is that this will allow us to computationally identify and emphasize aesthetic aspects of music. Specifically, we have identified an initial set of attributes (metrics) of music pieces on which to apply the Zipf–Mandelbrot law. These metrics include pitch of musical events, duration of musical events, the combination of pitch and duration of musical events, and several others. Our results are encouraging. We are working on automating our metrics so that we can test our hypothesis on a wide variety of music genres (baroque, classical, 20th century, blues, jazz, etc.). If this project is successful, we plan to investigate how such metrics may be used in various areas such as music education, music therapy, music recognition by computers, and computer–aided music analysis/composition.

1. Introduction

Computers have been used extensively in music to aid humans in analysis, composition, and performance [3, 4]. This is facilitated by the use of MIDI (Musical Instrument Digital Interface)—a coding scheme used to encode music data, such as pitch, duration and timbre [6].

Our project focuses on algorithmic techniques to help explore and identify aspects of beauty and balance in music. Specifically, we use MIDI renderings of various music genres (baroque, classical, 20th century, blues, jazz, etc.) to investigate the applicability of the Zipf–Mandelbrot law in analysis, composition, and performance of music via computer.

In this paper, we present results from analysis of 28 pieces from classical composers, namely Bach, Beethoven, Chopin, Debussy, Handel, and Schönberg. We compare these results to the MIDI rendering of a jazz standard and seven MIDI pieces containing randomly generated notes (random pitch, start time, and duration). Although preliminary, our results are encouraging. They strongly suggest that certain aspects of beauty and balance in music are algorithmically identifiable and classifiable.

This project has great potential to benefit music lovers by computationally identifying and emphasizing aesthetic aspects of music. Results from this project could lead to better human appreciation of music, including applications for music appreciation courses, and for exposing infants/young children to music. This project should also help develop improved techniques for recognition and classification of music by computers. Overall, this project could contribute significantly to areas such as music education, music therapy, music recognition by computers, and computer–aided music analysis/composition.

1.1. Zipf's Law

Zipf's law, named after Harvard University's Linguistics professor George Kingsley Zipf (1902–1950), is the observation that phenomena generated by self–adapting organisms, such as humans, follow *the principle of least effort* [1, 8, 15, 16].

Zipf discovered that if we plot the logarithm of the frequencies of all events in such a phenomenon against the logarithm of the rank of these events, we get a straight line with a slope of approximately –1.0. Figure 1 shows an example of a Zipf distribution discovered in Internet traffic.

In essence, Zipf's law states that, in any environment containing self–adapting agents able to interact with their surroundings, such agents tend to minimize their overall effort associated with this interaction (economy principle). That is, a system of such interacting agents tends to find a global optimum that minimizes overall effort [15]. This interaction involves some sort of exchange (e.g., information, energy, etc.). This agent interaction can be viewed as a phenomenon and every agent exchange as an event of this phenomenon.

These results can be extended to include more than two interacting agents—producers, consumers, or both producers and consumers. Zipf initially demonstrated his idea using

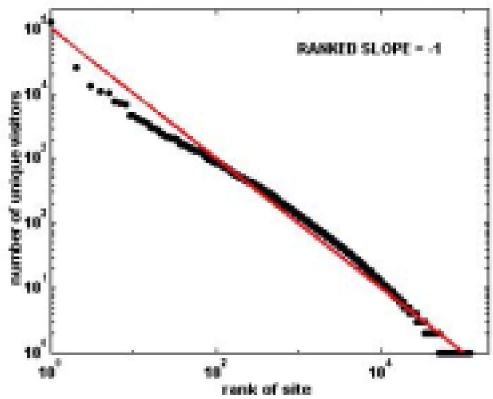

Fig. 1. Distribution of users among sites on the World-Wide-Web [2].

phenomena from natural language (e.g., words in a book). He successfully applied this principle to many other phenomena [15].

A similar theory was developed in the field of economics by Pareto (at the end of the 19^{th} century) [1, 12]. Zipf's (and Pareto's) work has (have) inspired and contributed to other fields studying the complexity of nature. Such fields include Fractals and Chaos Theory [10, 13]. In particular, Zipf's law was extended by Benoit Mandelbrot to account for similar distributions – phenomena that generate lines with slopes ranging between 0 (random phenomena) and negative infinity (monotonous phenomena). Such distributions are also called *power–law* distributions. Power–law distributions have been encountered in many man–made and naturally occurring phenomena including city sizes, incomes, word frequencies, earthquake magnitudes, thickness of sediment depositions, extinctions of species, traffic jams, and visits of websites [7].

Accordingly, researchers have discovered that certain phenomena (objects, artifacts) which follow Zipf's distribution are perceived by humans as pleasing, well–balanced, or even beautiful. For instance, the nautilus shell enhibits a near–Zipfian distribution in terms of its radii at 90-degree intervals (see figure 2).

In this light, Zipf's original law could be seen as describing phenomena that are ordered "just right" with respect to human sensory processes.

2. Music and the Zipf–Mandelbrot Law

Zipf mentions several occurrences of his distribution in musical pieces, including Mozart's 'Bassoon Concerto in Bb', Chopin's 'Etude in F–, Op. 25, No. 2', Irving Berlin's 'Doing What Comes Naturally', and Jerome Kern's 'Who' [15, pp. 336–337].

Voss and Clarke [14] measured several fluctuating physical variables, including output voltage of an audio amplifier, loudness fluctuations of music, and pitch fluctuations of music. Their samples included music from classical, jazz, blues, and rock radio stations collected continuously over 24 hours. Their results show that pitch and loudness fluctuations in music follow Zipf's distribution. However they were

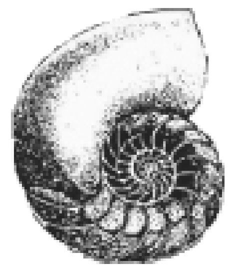

Fig. 2. Nautilus Shell

unable to show this for note fluctuations. Elliot and Atwell [5] also report that they did not find Zipf's distribution in note fluctuations. Both studies were carried out at the level of frequencies in an electrical signal.

Finally, Voss and Clark reversed the process so they could compose music through a computer. Their computer program used a Zipf's distribution (1/f power spectrum) generator to generate individual musical events (i.e., pitch and duration). The results were remarkable:

> *The music obtained by this method was judged by most listeners to be much more pleasing than that obtained using either a white noise source (which produced music that was 'too random') or a $1/f^2$ noise source (which produced music that was 'too correlated'). Indeed the sophistication of this '1/f music' (which was 'just right') extends far beyond what one might expect from such a simple algorithm, suggesting that a '1/f noise' (perhaps that in nerve membranes?) may have an essential role in the creative process.* [14, p. 258]

Nagai [11] extends this line of reasoning as follows:

> *The science of complexity locates life at the edge of chaos. Life is neither dead disorder nor dead order, but a critical point of phase–transition from chaos to order. In other words the life principle is neither white noise nor monotone, but a 1/f fluctuation.*
>
> *We are often "beside ourselves" with ecstasy over beautiful music or scenery, while we feel an aversion to high entropy such as a din or filth and weary of low entropy such as monotonous sound or figure, that is to say, we keep away from the two extremes, pure low entropy and utter high entropy. When you analyze the power spectrum of sound frequency, you will find that beautiful music shows 1/f frequency. You can also find 1/f fluctuations in curves rich in changes characteristic of beautiful natural scenery.* [11, p. 2]

3. Measurable Music Attributes

We have identified several attributes of music that could be used in deriving metrics (measurements) in search of Zipf–Mandelbrot distributions [9]. We suspect that some pieces may exhibit such distributions in one or more dimensions. These music attributes include the following: pitch, rests, duration, harmonic intervals (vertical), melodic intervals (horizontal), chords, movements, volume, timbre, tempo, dynamics. Some of these can be used independently, e.g., pitch; others can be used only in combinations, e.g., duration. Some lend themselves easily to the task, such as melodic intervals, whereas others not so easily, such as timbre. Of course, the fact that some music attributes may not seem good candidates for a Zipf–Mandelbrot distribution may be due to the fact that they have not traditionally been used as means of musical artistic expression, e.g., timbre.

Given our background as music listeners and musicians, we selected several of these music attributes and combinations of these music attributes to define particular metrics. These attributes were selected because (a) they have been used in earlier research, (b) have traditionally been used to express musical artistic expression and creativity, and/or (c) have been used in the analysis of composition. They are all studied extensively in music theory and composition. Obviously, this list of metrics is not complete. This is because there is probably no limit to the ways that sound could be used for artistic expression.

4. Implementation of Metrics

We are in the process of automating several of these metrics using Visual Basic. This will allow us to quickly test our hypothesis on hundreds of musical pieces readily available in MIDI rendering.

Some of the metrics have been challenging to implement. This is because music theoretic definitions are not always rigorous enough (e.g. harmonic intervals and melodic intervals). Thus, we attempted to define them in a more mathematically precise way, while attempting to preserve the "essence" of the listening experience for the corresponding music attribute.

Pitch: The number of times each of the 128 possible notes in a MIDI file occurs in a given piece of music.

Pitch Mod 12: The number of times each pitch of the 12–note chromatic scale occurs in a given piece of music.

Duration: The number of times that a note occurs at a specific duration, independent of the pitch of any given note.

Duration x Pitch: The number of occurrences of a particular duration of each of the 128 possible notes of the MIDI file. For example, it counts the number of occurrences of a middle C that has a duration of 12 time units, the number of occurrences of a middle C that has a duration of 15 time units, etc. These time units are not necessarily equivalent to standard music durations, due to possible performance (interpretation) variations.

Duration x Pitch Mod 12: The number of occurrences of a particular duration for each pitch of the 12-tone chromatic scale.

4.1. Interval–Based Metrics

These metrics consider various intervals that appear in a piece of music. A *melodic interval* is defined in terms of the change in pitch between successive notes over time. Experientially, such intervals correspond to a sense of "movement". A *harmonic interval*, on the other hand, is defined in terms of the difference in pitch between concurrent notes. Experientially, such intervals correspond to a sense of "color" or "mood".

Unfortunately, these definitions are not mathematically precise; for instance, arpeggiated chords produce chordal harmonies which are played over time. Such a chord would appear as a set of melodic intervals to an algorithm that treated changes in pitch over time as a melody, when it's actually harmonic in nature. In order to capture these events, we attempted to find a more rigorous definition, while at the same time capturing the basic essence (i.e., "movement" or "color") of each type of interval.

For our purposes, *harmonic interval* is an interval that occurs between two notes one of which is completely within the duration of the other. An interval between any two other notes is defined as a *melodic interval*. Specifically,

```
if ((S1 <= S2) and (E1 >= E2)) or
   (S1 >= S2) and (E1 <= E2) then
   interval is harmonic
else
   interval is melodic
```

where S_1 and E_1 are the starting and ending times of the first note and S_2 and E_2 are the starting and ending times of the second note.

We implemented each type of interval in a variety of ways. In some, we only counted certain intervals while discarding others, and in others we counted all interval occurrences in the entire piece. Due to limitations in space and given our experimental results, the following is only a partial list of the most promising metrics we have considered.

Our algorithm creates an array of notes from a MIDI file, sorts the notes by start time, and then sorts notes of the same start time in order of pitch. Below, we refer to *groups of notes* as notes sharing the same start time in this array.

Melodic Intervals: The number of times a particular pitch difference occurs between each note of a group and its closest note in the subsequent group.

All Melodic Intervals: The number of times a particular pitch difference occurs between each note of a group and each note of a group of the subsequent group.

Harmonic Intervals: The number of times a particular pitch difference occurs between the lowest note of a group and all other concurrent notes (see definition of harmonic interval).

All Harmonic Intervals: The number of times a particular pitch difference occurs between each note of a group and all other concurrent notes (see definition of harmonic interval).

Melodic and Harmonic Intervals: This metric is an amalgam of the 'Melodic Intervals' and 'Harmonic Intervals' metrics. It counts the number of times a particular pitch difference occurs between the lowest note in a group with every other note in the same group, and between each note in the group with the closest note in the subsequent group.

All Melodic and Harmonic Intervals: This metric is an amalgam of the ‘All Melodic Intervals’ and ‘All Harmonic Intervals’ metrics. It counts the number of times a particular pitch difference occurs between each note in a group and every note in the same group, and between each note in a group and every note in the subsequent group.

Higher Order Melodic Intervals: Given that melodic intervals capture *the change of pitches* over time, we attempt to capture higher orders of change. This includes the *changes in melodic intervals*, to the *changes of the changes in melodic intervals*, and so on. These higher–order metrics correspond to the notion of derivative in mathematics. Although a human listener may not be able to consciously hear such high–order changes in a piece of music, there may be some subconscious understanding taking place. In our research, we have found no mention of such intervals in music theory literature.

5. Experimental Study

We applied our metrics to a collection of quality MIDI renderings of music pieces in search of near–Zipfian distributions.

5.1. Music Pieces

The pieces used in our study are listed below by composer. When known, names of performers/arrangers appear in parentheses. Most MIDI renderings of classical pieces are from the Classical Archives <http://ftp.sunet.se/cma/>.

Bach

1. Two–Part Invention No. 13 in A–, BWV.784 (J.Sankey)
2. Chorale No. 229 'I stand before Thy Throne my God', BWV.350 (M. Yaskawa)
3. Chorale No. 6, BWV.656 (J.L.Dawson); Orchestral Suite No.3 in D '2. Air on the G String', BWV.1068 (M.Yaskawa)
4. Toccata and Fugue in D–, BWV.565
5. Jesu, Joy of Man's Desiring, BWV.147 (A.J.Krusz)
6. Portion of Toccata and Fugue in D–, BWV.565

Beethoven

1. Bagatelle No.25 in A– 'Fur Elise,' WoO.59 (B.Hisamori)
2. Romance No.1 in F, Op.40 (piano solo) (M. Yaskawa)

Chopin

1. Etude Op.10 No.12 in C– 'Revolutionary' (M.A.Deocariza).
2. Fantaisie–impromptu in C#–, Op.66.
3. Nocturne No.1 in Bb–, Op.9 No.1 (D.Inoue)

Debussy

1. Ballade (R.S.Finley)
2. Clair de lune (R.Sierra)
3. Première Rhapsodie pour Orchestre avec Clarinette (full orch.) (R.Bakels)
4. Rêverie (K.Stillwell)

Handel

1. Messiah – 17. Glory to God in the highest
2. Messiah – 2. Comfort ye, my people (D.L.Viens)
3. Messiah – 44. Hallelujah! For the Lord God omnipotent reigneth
4. The Musick for the Royal Fireworks (G.Pollen)

Mendelssohn

1. A Bee's Wedding (R.Finley)
2. Spinnerlied, Op.67 No.4 (C.Mote)

Schönberg

1. Pierrot Lunaire, Melodramas Op. 21 – Colombine
2. Pierrot Lunaire, Melodramas Op. 21 – Der Dandy
3. Pierrot Lunaire, Melodramas Op. 21 – Mondestrunken
4. Pierrot Lunaire, Melodramas Op. 21 – Valse de Chopin

Additionally, the following pieces were used for comparison purposes:

Rodgers & Hart

1. My Romance (Joe Pass)

Uniform–Distribution Random Number Generator

Six pieces (1–7) consisting of random note pitches, note start times, and note durations.

5.2. Results and Discussion

The following sections present average results for each metric. They also discuss what these numbers may mean with respect to our hypothesis.

Results are given in the format (slope, R^2). *Slope* is the slope of the trendline of the data values. R^2 is an indication of how closely the trendline fits the data values—the closer the fit, the more meaningful (reliable) the slope value.

Slope values range from 0 (high entropy—purely random) to $-\infty$ (low entropy—monotone). Slopes values near –1.0 correspond to Zipf’s distribution. R^2 range from 0.0 (extremely bad fit—data is all over the place) to 1.0 (perfect fit—data is already in a straight line).

Overall, these results support our hypothesis. Classical and jazz pieces averaged near-Zipfian distributions and tight fits across all metrics, whereas random pieces did not. Specifically, the across-metrics average slope for music pieces is –1.2653. The corresponding R^2 value, 0.8088, indicates a strong fit. The corresponding results for random pieces is –0.4763 and 0.6345, respectively.

Pitch

The classical pieces averaged (–1.3460, 0.7296) with a standard deviation of (0.2303, 0.0603). The values ranged from (–0.9680, 0.8596) to (–1.8701, 0.6076). Figure 3 shows the pitch distribution for one of these pieces. This is a Zipfian curve.

The jazz piece yielded (–1.3486, 0.8196).

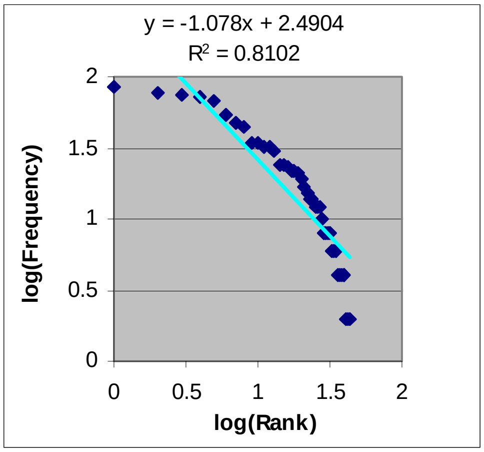

Fig. 3. Pitch distribution for Bach's Orchestral Suite No.3 in D '2. Air on the G String', BWV.1068.

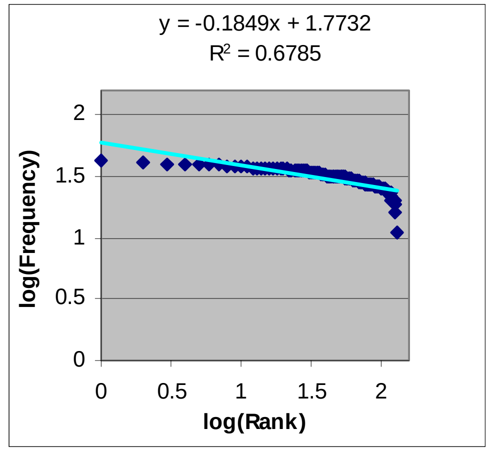

Fig. 4. Pitch distribution for Random Piece No. 7.

The random pieces averaged (–0.2060, 0.7027) with a standard deviation of (0.0606, 0.0426). The values ranged from (–0.1451, 0.7594) to (–0.3373, 0.6456). Figure 4 shows the pitch distribution for one of these pieces. This is not a Zipfian curve.

Pitch Mod 12

The classical pieces averaged (–1.0576, 0.6965) with a standard deviation of (0.5009, 0.1253). The values ranged from (–0.1549, 0.9579) to (–2.2230, 0.4455).

The jazz piece yielded (–1.0466, 0.7982).

The random pieces averaged (–0.0949, 0.8178) with a standard deviation of (0.0161, 0.1163). The values ranged from (–0.0736, 0.9797) to (–0.1256, 0.6555).

Remarks on Pitch Mod 12 Results

Interestingly, Schönberg's 12-tone pieces averaged (–0.2801, 0.7221) with a standard deviation of (–0.1549, 0.8762). The values ranged from (–0.2725, 0.9579) to (–0.4067, 0.9206). Obviously, this metric is very reliable in identifying 12-tone music, since it is characterized by the uniform distribution (high entropy) of pitches. The –0.2801 average agrees with most listeners' belief that 12-tone music is not very melodic.

The –2.2230 extreme was generated by Bach's Jesu, Joy of Man's Desiring, BWV.147—a piece that is relatively monotonous in terms of pitch (stays mostly in one key and has a characteristic theme repeated numerous times). This piece, however, is interesting (balanced, beautiful?) in terms of pitch durations and both harmonic and melodic intervals.

Finally, this metric may serve as a signature (unique identifier) for a piece. Such an identifier may be very useful in data retrieval applications, where one searches for different performances of a given piece among volumes of music. Specifically, this metric revealed that Bach's piece (6) was not a portion of Bach's piece (4) as identified by the Classical Archives, but actually the complete piece. Both pieces yielded a pitch-mod-12 result of (–1.0048, 0.7226). Results from the duration metric revealed that (4) and (6) were different performances. These observations were validated through listening and reading the corresponding MIDI scores.

Duration

The classical pieces averaged (–1.7880, 0.9101) with a standard deviation of (0.6966, 0.0499). The values ranged from (–0.9186, 0.9784) to (–3.9992, 0.7869).

The jazz piece yielded (–0.8429, 0.9314).

The random pieces averaged (–0.2304, 0.6766) with a standard deviation of (0.0809, 0.0596). The values ranged from (–0.1381, 0.7885) to (–0.3579, 0.6027).

Remarks on Duration Results

It should be noted that our data set was not uniformly quantized. Quantization works as a "snap-to-grid" mechanism for normalizing note durations in MIDI files. The quantizing software we used (Cakewalk Guitar Studio) did not have enough resolution (only up to 1/32 triplet). This affected the harmonic and melodic integrity of some fast pieces. For instance, in Bach's Toccata and Fugue in D– the available quantization reduced the melodic phrases in the opening of the piece to harmonic ones. The piece was virtually unrecognizable, as several neighboring tones were combined to form dissonant chords.

We are in the process of implementing our own quantizing software for normalizing our data set in terms of note durations up to 1/64 triplet. We feel that all duration-related results should be viewed cautiously and should be compared with results derived from a uniformly quantized data set.

Duration x Pitch

The classical pieces averaged (–0.8401, 0.8777) with a standard deviation of (0.3608, 0.0613). The values ranged from (–0.2514, 0.9684) to (–1.4140, 0.6995). Our average slope approximates Voss and Clarke's result derived from averaging 24 hours of classical music [14].

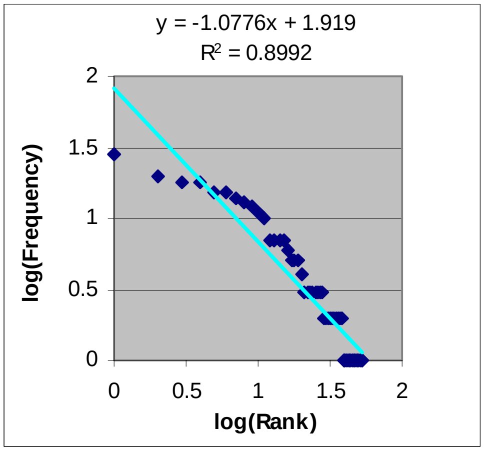

Fig. 5. Harmonic interval distribution for Bach's Two-Part Invention No. 13 in A–, BWV.784.

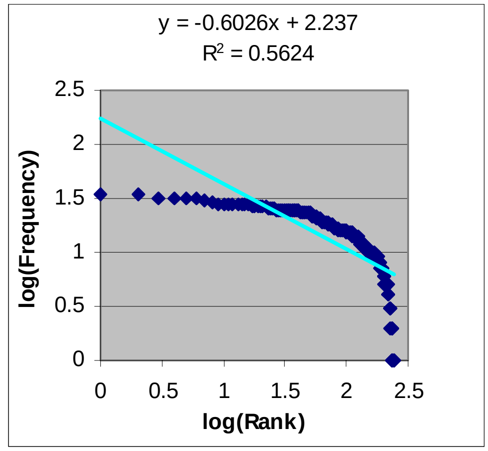

Fig. 6. Harmonic interval distribution for Random Piece No. 7.

The jazz piece yielded (–0.3118, 0.7488).

The random pieces averaged (–0.1545, 0.5620) with a standard deviation of (0.0722, 0.1554). The values ranged from (–0.0615, 0.7566) to (–0.2707, 0.3311).

Duration x Pitch Mod 12

The classical pieces averaged (–1.0655, 0.8926) with a standard deviation of (0.4242, 0.0500). The values ranged from (–0.3756, 0.9792) to (–1.8542, 0.7965).

The jazz piece yielded (–0.4084, 0.8390).

The random pieces averaged (–0.4368, 0.7497) with a standard deviation of (0.0753, 0.0518). The values ranged from (–0.329, 0.7959) to (–0.4978, 0.6461).

Melodic Intervals

The classical pieces averaged (–1.4336, 0.8590) with a standard deviation of (0.2072, 0.0647). The values ranged from (–0.9979, 0.9313) to (–1.9430, 0.6757).

The jazz piece yielded (–1.1430, 0.7719).

The random pieces averaged (–-0.7073, 0.6014) with a standard deviation of (0.0559, 0.0449). The values ranged from (–0.6228, 0.6850) to (–0.7968, 0.5430).

All Melodic Intervals

The classical pieces averaged (–1.3700, 0.7820) with a standard deviation of (0.1834, 0.0738). The values ranged from (–0.9704, 0.9243) to (–1.7038, 0.6487).

The jazz piece yielded (–1.1181, 0.7452).

The random pieces averaged (–0.6486, 0.6281) with a standard deviation of (0.0448, 0.0704). The values ranged from (–0.5887, 0.7532) to (–0.7281, 0.5287).

Harmonic Intervals

The classical pieces averaged (–1.2178, 0.8229) with a standard deviation of (0.2244, 0.0583). The values ranged from (–0.8951, 0.9058) to (–1.8011, 0.7167). Figure 5 shows the harmonic interval distribution for one of these pieces. This is a Zipfian curve.

The jazz piece yielded (–1.0815, 0.7963).

The random pieces averaged (–0.6350, 0.5965) with a standard deviation of (0.0454, 0.0663). The values ranged from (–0.5857, 0.7379) to (–0.7028, 0.5381). Figure 6 shows the harmonic interval distribution for one of these pieces. This is not a Zipfian curve.

All Harmonic Intervals

The classical pieces averaged (–1.3247, 0.8179) with a standard deviation of (0.2736, 0.0562). The values ranged from (–0.9496, 0.9018) to (–2.0076, 0.7307).

The jazz piece yielded (–1.1257, 0.8110).

The random pieces averaged (–0.6652, 0.5462) with a standard deviation of (0.0274, 0.0615). The values ranged from (–0.6066, 0.6767) to (–0.6896, 0.4903).

Melodic and Harmonic Intervals

The classical pieces averaged (–1.3894, 0.7900) with a standard deviation of (0.1863, 0.0615). The values ranged from (–1.0720, 0.9041), to (–1.8610, 0.6383).

The jazz piece yielded (–1.0781, 0.7314).

The random pieces averaged (–0.6963, 0.5517) with a standard deviation of (0.0381, 0.0512). The values ranged from (–0.6404, 0.6589) to (–0.7585, 0.5114).

All Melodic and Harmonic Intervals

The classical pieces averaged (–1.3858, 0.7371) with a standard deviation of (0.1894, 0.0565). The values ranged from (–0.9909, 0.8215), to (–1.8017, 0.6508).

The jazz piece yielded (–1.1163, 0.6933).

The random pieces averaged (–0.6765, 0.5246) with a standard deviation of (0.0346, 0.0620). The values ranged from (–0.6208, 0.6497) to (–0.7391, 0.4671).

6. Conclusion

Our project explores different ways to apply the Zipf-Mandelbrot law on musical pieces encoded in MIDI. Our hypothesis is that this will allow us to computationally identify and perhaps emphasize aesthetic aspects of music. Our results, although preliminary, support this hypothesis. Specifically, the across-metrics average slope for music pieces is near-Zipfian (–1.2653). The corresponding R^2 value indicates a strong fit (0.8088). The corresponding results for random pieces is –0.4763 and 0.6345, respectively.

These results agree with the observation that all music pieces (as opposed to the random ones) in our data set are well-known, celebrated pieces of music. It is reasonable to assume that their popularity is correlated to the fact that listeners find them balanced, pleasant, perhaps harmonious, or even beautiful.

It is not clear if the pattern observed in these results will generalize, i.e., whether it holds for all music, or it is coincidental. Obviously, applying this metric to a wide variety of pieces will provide an empirical answer to this question. We intend to test our hypothesis on a wide variety of composers and music genres (e.g., baroque, classical, 20^{th} century, blues, and jazz).

6.1. Future Work

We expect that our metrics, regardless of whether or not they support our hypothesis, may provide a "signature" mechanism—a way to identify music pieces and perhaps help to automatically classify them in terms of composer or genre. Such classification may be implemented through connectionist means. A neural network could be trained on arithmetic results (e.g., slope of a trendline in a graph) derived from applying our metrics to a wide variety of musical pieces. Such a neural network could be applied in music identification and/or classification tasks.

These and other metrics under development could easily be incorporated into a collection of tools for computer–aided music composition. Such tools may be analytical or generative in nature.

In terms of analysis, such tools could be used by music composers for formative evaluation—that is to measure the entropy/balance (beauty?) of a musical piece under development. Examples of such feedback include whether a piece is too monotone or too random, and whether it is similar to, say, Bach's Toccata and Fugue in D– along certain music–theoretic dimensions. Obviously such metrics should be used with caution, similarly to readability metrics in natural language composition. Such metrics are, at best, crude abstractions which, if followed blindly, could limit a composer's creativity and range of expression. On the other hand, if applied properly, they could enhance/facilitate artistic expression.

In terms of music generation, such metrics could be used in computer–aided music composition. Specifically, using Artificial Intelligence techniques (e.g., genetic algorithms), one could perhaps reverse the process incorporated in the above metrics and generate computer music that sounds "pleasing, beautiful, harmonious." This will be in line with Voss and Clarke's music generation experiment mentioned in section 2. A composer could perhaps contribute an initial seed (a melodic outline) and have the system expand it automatically. Alternatively, tools could be developed for allowing a composer to supply a music phrase for automatic generation of variations. This would be similar to various visual effects (transformations) available in all mainstream graphic packages for manipulation/enhancement of photographs.

Acknowledgements

The authors would like to thank James Wilkinson and Robert Aiken for their suggestions and comments, and Valerie Sessions for her help during the early phases of this work.

References

1. Adamic, L.A., (1999), "Zipf, Power–laws, and Pareto – a Ranking Tutorial", www.parc.xerox.com/istl/groups/iea/papers/ranking/
2. Adamic, L.A., and Huberman, B.A. (2000), "The Nature of Markets in the World Wide Web", *Quarterly Journal of Electronic Commerce*, **1(1)**:5–12.
3. Balaban, M., Ebcioğlu, K., and Laske, O., eds. (1992), Dobrian, C. (1988), *Understanding Music with AI: Perspectives on Music Cognition*, AAAI Press and MIT Press.
4. Dobrian, C. (1992), "Music and Artificial Intelligence", www.arts.uci.edu/dobrian/CD.music.ai.htm
5. Elliot, J. and Atwell, E. (2000), "Is Anybody Out There? The Detection of Intelligent and Generic Language–Like Features", *Journal of the British Interplanetary Society 53(1/2)*, pp. 13–22, www.comp.leeds.ac.uk/eric/jbisjournal2000.ps
6. Glatt, J., "Tutorial for MIDI Users", www.borg.com/~jglatt/tutr/miditutr.htm
7. Knuth, K. (1997), "Power Laws and Hierarchical Organization in Complex Systems—From Sandpiles and Monetary Systems to Brains, Language, and Music", *CUNY Cognitive Science Symposium*. http://bulky.aecom.yu.edu/users/kknuth/complex/powerlaws.html .
8. Li, W. (2000), "Zipf's Law" http://linkage.rockefeller.edu/wli/zipf/
9. Manaris, B., McCormick, C., and Purewal, T. (2001), "Searching for Beauty in Music–Applications of Zipf's Law in MIDI-Encoded Music", 2001 Sigma Xi Forum, *Science, the Arts and the Humanities: Connections and Collisions*, Raleigh, NC.
10. Mandelbrot, B.B. (1977), *The Fractal Geometry of Nature*, W.H. Freeman and Company.

11. Nagai, T. (2000), “The Rhythm of Life”, www.themestream.com/articles/133106.html
12. Pareto, V. (1897), *Cours d'Economie Politique*, Rouge (Lausanne et Paris). Cited in [10].
13. Schroeder, M. (1991), *Fractals, Chaos, Power Laws*, W.H. Freeman.
14. Voss, R.F., and Clarke, J. (1975), “1/f Noise in Music and Speech”, *Nature* 258, pp. 317–318.
15. Zipf, G.K. (1949), *Human Behavior and the Principle of Least Effort*, Addison–Wesley.
16. Zipf, G.K. (1965), *Psycho–Biology of Languages*, MIT Press (original publication Houghton–Mifflin, 1935).