Can Beautiful Music be Recognized by Computers?

Nature, Music, and the Zipf–Mandelbrot Law

Bill Manaris, Charles McCormick, and Tarsem Purewal
Computer Science Department
College of Charleston

Technical Report CoC/CS TR#2002-7-1
March 1, 2002

Keywords: Music Analysis, Zipf–Mandelbrot Law, Golden Ratio, MIDI, Artificial Intelligence.

Abstract

We discuss the application of the Zipf–Mandelbrot law to musical pieces encoded in MIDI. Specifically, we have identified a set of metrics on which to apply this 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 hypothesis is that these metrics will allow us to computationally identify and perhaps emphasize aesthetic aspects of music. We have developed a system that automates calculation of these metrics in MIDI-encoded pieces. Using this tool we conducted a study on a corpus of 220 pieces from baroque, classical, romantic, 12-tone, jazz, rock, DNA strings, and random music. Our results support the above hypothesis. We discuss limitations, and how to minimize the potential for statistical error through composite metrics. We present possible applications and an overview of preliminary work in computer–aided music composition.

1. Introduction

Webster's New World Dictionary defines beauty as *the quality attributed to whatever pleases or satisfies the senses or mind, as by line, color, form, texture, proportion, rhythmic motion, tone, etc., or by behavior, attitude, etc.* One of the difficulties in addressing the question posed in the title (and in pursuing this research) is that there exists no precise, quantitative definition for beauty.

Since computers are ultimately manipulators of quantitative representations, any attempt to model non-quantifiable qualities is inherently problematic.[1] In the case of beauty, an additional problem is that it appears to be affected by subjective (cultural, educational, physical) biases of an individual – that is, beauty is in the eye (ear, etc.) of the beholder. Or is it?

[1] This is a traditional, well-founded concern about the efficacy of Artificial Intelligence (AI) research. Nevertheless, AI research has generated many valuable applications to date.

Research in fractals and chaos theory suggests that aspects of our perception of beauty, balance, and harmony in nature may be guided by inherent, subconscious processes [12, 19]. Specifically, it appears that the design and assembly of man-made or natural objects (complex systems) is guided by hidden rules that impose constraints on how structures are put together [7, 8, 18].

> *All of these rules can be found – once one knows how to look for them – in traditional art and architecture. The rules guarantee a connection to the human observer, who notices (albeit subconsciously) the mathematical ordering inherent in a pleasing design. For example, design units are very rarely single; they usually recur with some multiplicity. Another way of saying this is that [our inherent] need for ordering leads us to make certain subunits similar to each other. The amazing thing is that nature also creates structural subunits that are similar.* [18, p. 910]

In the context of music, Voss and Clarke [21, 22] have suggested that music might also be viewed as a complex system whose design and assembly is partially guided by rules subconscious to the composer. They have also demonstrated that listeners may be guided by similar rules in their aesthetic response to a music piece (see section 4).

Our project focuses on algorithmic techniques to help explore and identify aspects of beauty and balance in music. Specifically, we have used MIDI renderings of various musical genres (baroque, classical, 20th century, rock, jazz, etc.) to investigate the applicability of the Zipf–Mandelbrot law in analysis and composition of music via computer. Results from analysis of 220 pieces suggest that certain aspects of beauty in music are algorithmically identifiable and classifiable. We also discuss an application of these results in computer-aided music composition.

2. The Zipf–Mandelbrot Law

Zipf's law is named after George Kingsley Zipf (1902–1950), a linguistics professor at Harvard University. In a highly influential book [23], Zipf proposed that human behavior might be studied as a natural phenomenon like everything else in nature.

He also observed that phenomena generated by complex social or natural systems (such as human language and music), tend to follow a statistically predictable structure. This structure may be present at various hierarchical levels [8]. Specifically, if we plot the frequencies of words in a book, such as Homer's Iliad, against their statistical rank on logarithmic scale, we get a straight line with a slope of approximately –1.0. In other words, the probability of occurrence of words starts high and decreases rapidly. A few words, such as 'a' and 'the', occur very often, whereas most words, such as 'unconditional', occur rarely. Formally, the frequency of occurrence of the n^{th} ranked word is $1/n^a$, where a is close to 1 [9].

In essence, Zipf's law states that, in any system consisting of interacting entities, the overall effort associated with this interaction is minimized (economy principle) [23, 24]. This interaction involves some sort of exchange (e.g., information, energy, meaning, etc.). Examples include words in human languages, computer languages, operating system calls, colors in images, city sizes, incomes, word frequencies, music, earthquake magnitudes, thickness of sediment depositions, extinctions of species, traffic jams, and visits of websites [1, 2, 8, 9, 12, 19]. For instance, figure 1 shows an example of a Zipf distribution exhibited by Internet traffic.

Similar laws have been developed independently by Pareto, Lotka, and Bendford [1, 9, 12]. These laws have inspired and contributed to other fields studying the complexity of nature. In particular, Zipf's law inspired and was extended by Benoit Mandelbrot to account for a wider range of natural phenomena. Such phenomena may generate lines with slopes ranging between 0 (random phenomena) and negative infinity (monotonous phenomena). These distributions are also known as *power–law* distributions [8, 19].

3. The Golden Ratio

The ancient Greeks, at least as early as Pythagoras (569-475BC), knew of the *golden ratio* (aka golden section or golden mean). According to Pythagoras,

> *The Golden Ratio manifests in the whole of creation. Take the ratio of the length of a man and the height of his navel. The ratio of the sides of the Great Temple. The ratio between the long and short sides of a pentagram. Why is this? Because the ratio of the Whole to the Greater is the ratio of the Greater to the Lesser.*

The golden ratio is the irrational number (1.618034…). It is often represented by the Greek letter Phi (φ). A closely related value is just the decimal part of Phi (0.618034…).

The golden ratio is found in beautiful natural objects, such as the nautilus shell (see figure 2). Phi is also found in aesthetically pleasing human artifacts. This includes the shapes of ancient buildings, such as the Parthenon and Egyptian pyramids. It also includes beautiful works of famous artists, such as da Vinci and Michelangelo [7].

Researchers have found that many composers, including Mozart, Bartók, Debussy, Schubert, Bach, and Beethoven have used the golden ratio in their compositions [7, 6, 13, 14, 15].

Perhaps not surprisingly, objects and artifacts incorporating the golden ratio are also describable under the Zipf-Mandelbrot law. For instance, in the nautilus shell, the radii at 90 degree-intervals approximate Zipf's distribution (see figure 2). Accordingly, researchers have observed that objects and artifacts that follow Zipf's distribution are often perceived as pleasing, well–balanced, or even beautiful. In this light, Zipf's law could perhaps be seen as describing phenomena that are ordered "just right" with respect to human sensory processes.

4. 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' [23, pp. 336–337].

Voss and Clarke [21] 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

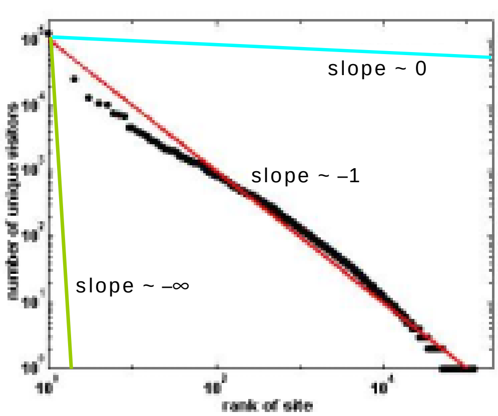

Fig. 1. Distribution of users among sites on the World-Wide-Web [2].

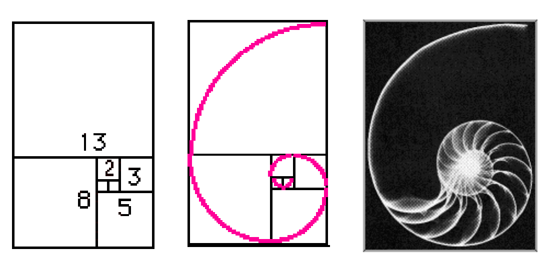

Fig. 2. The golden ratio and Nautilus shell [7].

fluctuations in music follow Zipf's distribution. However they were 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 [21, 22]. Their computer program used a Zipf's distribution (1/f or pink-noise) 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.* [21, p. 318]

5. Measurable Musical Attributes

We have identified several attributes of musical that could be used in deriving metrics in search of Zipf–Mandelbrot distributions [10, 11]. We suspect that some pieces may exhibit such distributions in one or more dimensions. These musical 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 musical 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 musical attributes and combinations of these musical attributes to define particular metrics. These attributes were selected because they (a) 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.

6. Implementation of Metrics

We have automated several of these metrics using Visual Basic and C++. This allowed us to quickly test our hypothesis on hundreds of musical pieces readily available in MIDI format. The following is a brief overview of these metrics.

6.1. Metrics Based on Individual Notes

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.

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

6.2. Metrics Based on Intervals

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 system 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.

Harmonic Intervals: The number of times a particular pitch difference occurs between the lowest note of a group and all other concurrent notes.

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.

Harmonic Bigrams: The number of times a particular pair of harmonic intervals occurs in a piece. This metric captures the balance of harmonic-triad (chord) structures.

Melodic Bigrams: The number of times a particular pair of melodic intervals occurs in a piece. This metric captures the balance of melodic structure.

Melodic Trigrams: The number of times a particular triplet of melodic intervals occurs in a piece. This metric also captures the balance of melodic structure.

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.

7. Experimental Study

We calculated our metrics on a collection of quality MIDI renderings of musical pieces in search of near–Zipfian distributions. Additionally, we included a set of DNA-generated pieces and a set of random pieces for comparison purposes. Most MIDI renderings of classical pieces are from the Classical Archives [20].

7.1. Musical Pieces

Our corpus consisted of 220 MIDI pieces. Due to space limitations, we summarize them below by genre and composer.

Baroque: Bach, Buxtehude, Corelli, Handel, Purcell, Telemann, and Vivaldi (38 pieces).

Classical: Beethoven, Haydn, and Mozart (18 pieces).

Early Romantic: Hummel, Rossini, and Schubert (14 pieces).

Romantic: Chopin, Mendelssohn, Tarrega, Verdi, and Wagner (29 pieces).

Late Romantic: Mussorgsky, Saint-Saens, and Tchaikovsky (13 pieces).

Post Romantic: Dvorák and Rimsky-Korsakov (13 pieces).

Modern Romantic: Rachmaninov (2 pieces).

Impressionist: Ravel (1 piece).

Modern (12 Tone): Berg, Schönberg, and Webern (15 pieces).

Jazz: Charlie Parker, Chick Corea, Cole Porter, Dizzy Gillespie, Django Reinhardt, Duke Ellington, John Coltrane, Miles Davis, Sonny Rollins, and Thelonius Monk (33 pieces).

Rock: Black Sabbath, Led Zeppelin, and Nirvana (12 pieces).

Pop: Beatles, Bee Gees, Madonna, Mamas and Papas, Michael Jackson, and Spice Girls (18 pieces).

Punk Rock: The Ramones (3 pieces).

DNA Encoded Music: actual DNA sequences encoded into MIDI, and simulated DNA sequences (12 pieces).

Random (White Noise): music consisting of random note pitches, note start times, and note durations "composed" by a uniformly-distributed random number generator (6 pieces).

Random (Pink Noise): music consisting of random note pitches, note start times, and note durations "composed" by a random number generator exhibiting Zipf's distribution (6 pieces).

7.2. Results

Due to space limitations, we show only average results for each genre in terms of slope, R^2, and corresponding standard deviations (see Table 1).

Slope is the slope of the trendline of the data values. Slopes may range from 0 (high entropy—purely random) to $-\infty$ (low entropy—monotone). Slopes near –1.0 correspond to Zipf's distribution.

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. R^2 may range from 0.0 (extremely bad fit—data is all over the graph) to 1.0 (perfect fit—data is already in a straight line). We considered R^2 values larger than 0.7 to be a good fit.

Genre	Slope	R^2	Slope Std	R^2 Std
Baroque	-1.1784	0.8114	0.2688	0.0679
Classical	-1.2639	0.8357	0.1915	0.0526
Early Romantic	-1.3299	0.8215	0.2006	0.0551
Romantic	-1.2107	0.8168	0.2951	0.0609
Late Romantic	-1.1892	0.8443	0.2613	0.0667
Post Romantic	-1.2387	0.8295	0.1577	0.0550
Modern Romantic	-1.3528	0.8594	0.0818	0.0294
Impressionist	-0.9186	0.8372	*N/A*	*N/A*
12 Tone	-0.8193	0.7887	0.2461	0.0964
Jazz	-1.0510	0.7864	0.2119	0.0796
Rock	-1.2780	0.8168	0.2967	0.0844
Pop	-1.2689	0.8194	0.2441	0.0645
Punk Rock	-1.5288	0.8356	0.5719	0.0954
DNA	-0.7126	0.7158	0.2657	0.1617
Random (Pink)	-0.8714	0.8264	0.3077	0.0852
Random (White)	-0.4430	0.6297	0.2036	0.1184

Table 1. Average results across metrics for each genre.

7.3. Discussion

Overall, the results indicate that aspects of beauty in music may be algorithmically identifiable and classifiable. The average for all musical pieces (excluding DNA, pink random, and white random pieces) across all metrics is –1.2004, a near-Zipfian distribution; the corresponding fit across all metrics is 0.8213.

For illustration purposes, figure 3 shows the pitch distribution for Bach's Orchestral Suite No.3 in D 'Air on the G String', BWV.1068. Its slope is –1.078, and its fit is 0.8102. Figure 4 shows the harmonic interval distribution for Bach's Two-Part Invention No. 13 in A minor, BWV.784. Its slope is –1.0776, and its fit is 0.8992.

The reader should compare these to examples from random piece No. 7. This piece was "composed" via a uniform-distribution (white noise) random number generator. Figures 5 and 6 show the pitch distribution and harmonic-interval distributions, respectively, for this piece. The corresponding slopes are –0.1849 and –0.6026, respectively.

Remarks on Pitch Mod 12 Results

Study of individual metrics reveals several patterns perhaps beyond the scope of this paper, but too interesting to not mention. One such pattern is related to the Pitch-Mod-12 metric. As mentioned earlier, this metric captures the number of times each pitch of the 12–note chromatic scale occurs in a given piece of music.

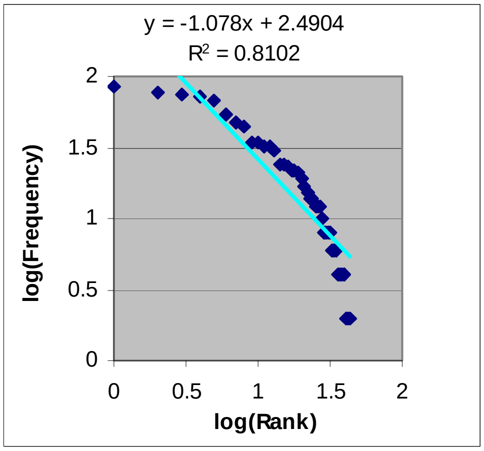

Fig. 3. Pitch distribution for Bach's Orchestral Suite No.3 in D 'Air on the G String', BWV.1068.

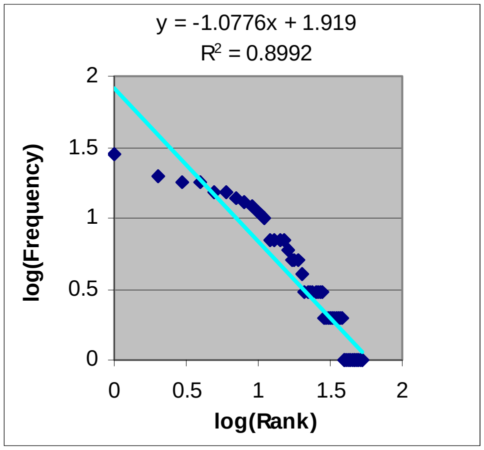

Fig. 4. Harmonic interval distribution for Bach's Two-Part Invention No. 13 in A–, BWV.784.

In the case of 12-tone music, we expected to see slopes suggesting uniform distribution (slope close to 0). Indeed, the corresponding slope for 12-tone pieces averaged –0.3168 with a standard deviation of 0.1801. In particular, Schönberg's pieces averaged –0.2801 with a standard deviation of –0.1549. This was comparable to the average for random (white noise) pieces, namely –0.1535. Obviously, this metric is very reliable in identifying 12-tone music, since such metric is characterized by the uniform distribution of pitches.

For comparison purposes, the next closest average slope for musical pieces was exhibited by Jazz pieces (–0.8770), followed by Late Romantic ones (–1.0741).

8. Limitations

Statistical approaches, such as the one reported herein, are abstractions. If successful, they tend to summarize the essential while abstracting (hiding) the unessential. But this is not always so. Nettheim [16] describes this limitation of statistical approaches in musicology as follows:

Statistics deals largely with replication, whereas music deals with particular cases … A particular case, such as an outlier, may occasionally interest a statistician who is otherwise concerned with general tendencies, but in music particular cases are everything. Changing a single note by the smallest amount (say a C to C sharp) may have little statistical but enormous musical effect, for by its nature a musical masterpiece is an organic whole, not just a series of note-decisions. [16, p. 94]

In our case, for any given metric, one could construct many "pathological" examples – MIDI renderings that the metric would identify as balanced, which are nevertheless unbalanced. For instance, one could take Bach's Toccata and Fugue in D minor, whose Pitch-Mod-12 slope is –1.0048, and reorder the notes in increasing order of MIDI Pitch-Mod-12 value. The resultant piece would then begin with one long sequence of all C notes grouped together, followed by all C sharp notes, and so on. This piece would sound very unbalanced (to say the least), but still exhibit a Pitch-Mod-12 slope of –1.0048.

From a practical perspective, this limitation may be addressed though composite metrics which take hierarchical structure into account. Our research indicates that combining metrics into a weighted composite for a single piece reduces the possibility for such error (see section 9).

Nevertheless, the existence of pathological examples such as the one described above indicates that **a near-Zipfian distribution is a necessary, but not sufficient condition for beautiful music**.

Another limitation of the study is that it focused only on the level of whole pieces (as opposed to individual movements, phrases, etc.). However, beauty may be encountered at various levels of granularity (e.g., piece, movement, phrase). Obviously, our metrics may easily be applied at various levels of granularity – as long as the sample can be encoded in MIDI. Based on observations in other complex systems [8], we expect that after a certain threshold in terms of sample size (too small), the slopes observed in musical excerpts will no longer exhibit a near-Zipfian distribution and perhaps may converge to some other slope.

Finally, we assumed that the MIDI archives we consulted over the Internet provided appropriate raw material for our study. In particular, Nettheim questioned Voss' accumulation of 24-hour continuous samples from radio stations because it did not focus on single pieces, which are normally the largest unit of artistic significance. This may have possibly introduced statistical correlations of unclear musical significance [15].

In our case, some of the samples we studied included movements diverse in terms of tonality, tempo, and mood, whereas others did not. Also, some of the samples were not complete pieces; instead they contained only certain movements from a larger piece. Although this may have also introduced inappropriate statistical correlations, one could argue that the selected samples were included in the

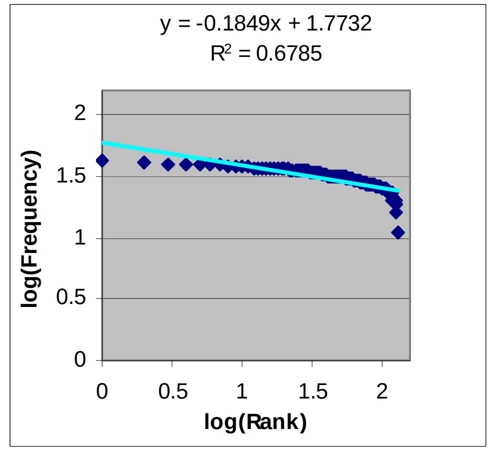

Fig. 5. Pitch distribution for Random (White Noise) Piece No. 7.

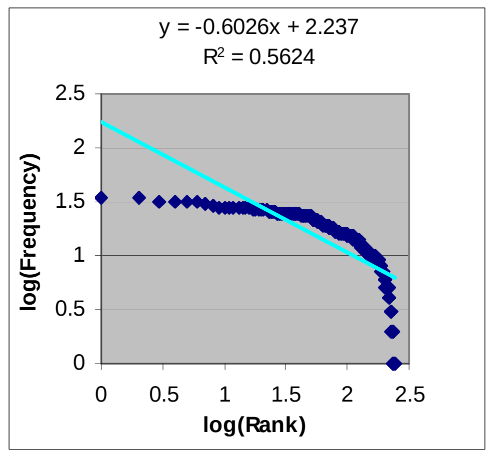

Fig. 6. Harmonic interval distribution for Random (White Noise) Piece No. 7.

MIDI archives because they were viewed as complete units of artistic significance.

9. Composite Metrics

Our study suggests that combining metrics into a weighted composite (consisting of metrics that capture various aspects throughout the possible space of measurable aesthetic attributes) minimizes the possibility for statistical error.

Specifically, each metric captures a different aspect of balance within a piece. For instance, melodic interval metrics capture the balance of melodic movement in a piece. Harmonic interval metrics capture the balance of harmonic movement in a piece. A duration-based metric may capture the beauty in the interpretation of a particular piece (i.e., the difference between quantized note and actual note durations). Bigram and trigram metrics capture aspects of structural beauty of a piece. And so on.

We have experimented with composite metrics having (a) various weights assigned to individual metrics and (b) conditional combinations of individual metrics. Such composite metrics appear promising. For instance, they could be used to identify pieces that have similar aesthetic characteristics to a given piece.

Additionally composite metrics may help derive a statistical signature (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. For instance, during an earlier study [11], we discovered a mislabeled MIDI piece by noticing that it had identical Pitch-Mod-12 slope and R^2 values with another MIDI piece. The two files contained different interpretations (performances) of Bach's Toccata and Fugue in D minor.

10. Conclusion

Our study indicates that it is possible to computationally identify aesthetic aspects of music. Our approach may be used to computationally identify and, in some cases, classify pieces based on their Zipfian distributions (or lack thereof) across a variety of metrics. Moreover, our results suggest that a near-Zipfian distribution is a necessary (but not sufficient) condition for beautiful music.

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.

10.1. Music Analysis

In terms of analysis, such tools could be used by composers for formative evaluation—that is to measure the 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 minor 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.

10.2. Music Generation

In terms of music generation, we are currently exploring how the above metrics may be used to guide a computer-aided music composition program.

Specifically, we have replicated Voss and Clarke's pink-noise generation approach [21, 22] to create music exhibiting a Zipfian distribution across many of our metrics. However, this music produced near-zero distributions with respect to bigram- and trigram-based metrics. This is because it lacks harmonic and melodic structure.

Our approach incorporates traditional Artificial Intelligence techniques (e.g., genetic algorithms and search) to add structure to such music [3, 4]. Specifically, a composer contributes an initial seed (a musical phrase) and a melodic outline. The system expands the seed automatically guided by the melodic outline. Our objective is to use the slope and R^2 values for various metrics as guides in determining a suitable organization of parts into a musical structure.

The musical structure is specified as a YACC ("Yet Another Compiler Complier") grammar. This allows a top-down, hierarchical description of a piece. For example, a piece of music may be broken down into a collection of themes. These themes can be further broken down into collections of musical phrases, which can in turn be broken down into collections of musical motives. YACC provides a mechanism for characterizing these divisions as subsets or supersets of each other. The largest defined structure will be the complete musical piece, and the smallest defined structure will be the initial input into the program.

By anticipating future tokens through probabilistic means, we have implemented a probabilistic pushdown automaton that guides the generation process. YACC actions are used to enforce constraints and long-term dependencies, and evaluate resultant substructures in terms of the metrics described above.

Alternatively, tools could be developed for allowing a composer to supply a musical 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 Valerie Sessions for their help during the early phases of this work. They also acknowledge the support of the College of Charleston through an internal R&D grant.

References

1. Adamic, L.A., (1999), "Zipf, Power–laws, and Pareto – a Ranking Tutorial", http://ginger.hpl.hp.com/shl/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), pp. 5–12.
3. Balaban, M., Ebcioğlu, K., and Laske, O., eds. (1992), *Understanding Music with AI: Perspectives on Music Cognition*, Cambridge: AAAI Press/MIT Press.
4. Cope, D. (2001), *Virtual Music*, Cambridge: MIT Press.
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. Howat, R. (1983). *Debussy in Proportion*. Cambridge: Cambridge University Press (cited in [16]).

7. Knott, R. (2002), "Fibonacci Numbers and Nature", www.mcs.surrey.ac.uk/Personal/R.Knott/Fibonacci/fibnat.html .

8. 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 .

9. Li, W. (2002), "Zipf's Law", http://linkage.rockefeller.edu/wli/zipf/ .

10. 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.

11. Manaris, B., Purewal, T., and McCormick, C. (2002), "Progress Towards Recognizing and Classifying Beautiful Music with Computers", 2002 IEEE SoutheastCon, Columbia, SC (to appear).

12. Mandelbrot, B.B. (1977), *The Fractal Geometry of Nature*, New York: W.H. Freeman.

13. Marillier, C.G. (1983). "Computer Assisted Analysis of Tonal Structure in the Classical Symphony", *Haydn Yearbook 14*, pp. 187-199 (cited in [16]).

14. May, M. (1996), "Did Mozart Use the Golden Section?", *American Scientist 84*(2). www.sigmaxi.org/amsci/issues/Sciobs96/Sciobs96-03MM.html .

15. Nettheim, N. (1992), "On the Spectral Analysis of Melody", *Interface: Journal of New Music Research 21*, pp. 135-148.

16. Nettheim, N. (1997), "A Bibliography of Statistical Applications in Musicology", *Musicology Australia 20*, pp. 94-106.

17. Pareto, V. (1897), *Cours d'Economie Politique*, Rouge (Lausanne et Paris). Cited in [12].

18. Salingaros, N..A., and West, B.J. (1999), "A Universal Rule for the Distribution of Sizes", *Environment and Planning B*(26), pp. 909-923.

19. Schroeder, M. (1991), *Fractals, Chaos, Power Laws*, New York: W.H. Freeman.

20. The Classical Music Archives, www.classicalarchives.com.

21. Voss, R.F., and Clarke, J. (1975), "1/f Noise in Music and Speech", *Nature* 258, pp. 317–318.

22. Voss, R.F., and Clarke, J. (1978), "1/f Noise in Music: Music from 1/f Noise", *Journal of Acoustical Society of America 63*(1), pp. 258–263.

23. Zipf, G.K. (1972), *Human Behavior and the Principle of Least Effort*, New York: Addison–Wesley (original publication Hafner Publishing Company, 1949).

24. Zipf, G.K. (1965), *Psycho–Biology of Languages*, Cambridge: MIT Press (original publication Houghton–Mifflin, 1935).