

Perspectives and Interpretations: Weighted Power Law Analyses of Three Short Works for Piano

DOUGLAS SCOTT

*Junior Research Fellow: Manfred and Penny Conrad Institute for Music Therapy Research
Wilfrid Laurier University*

KEYWORDS

Music, Power laws, Zipf's law, Noise profiles, Timescape, Wavelet, Schumann, Chopin, Mendelssohn

0. ABSTRACT

Power law relationships, which describe scaling relationships of data, are powerful information theoretic descriptive tools in many empirical contexts, including that of music. Zipf's law (pink noise) describes the optimum case of a power law relation where observations are exactly inversely proportional to their rank. Descriptions that approximate a pink noise signature can be said to maximize the amount of information in a signal, and is thus suggestive of a richness of an understanding. This information density of pink noise signatures is not, however, necessarily a desirable quality for explanations in general, which, by definition, "flattens out" some data while highlighting others, ideally those most relevant to an interpretation. The privileging of data most relevant to comprehension corresponds to a red noise relationship, as opposed to the pink noise of Zipf's law or the white noise of a description that highlights nothing in particular. Here, I explore and evaluate this concept of red noise explanations in the form of analyses of three comparable short piano works: Robert Schumann's "Von fremden Ländern und Menschen" (no.1 from Kinderszenen op. 15, 1838), Frédéric Chopin's Prelude op. 28 no. 20 (1838-9), and Felix Mendelssohn's "Venetianisches Gondellied" (no. 6 from Lieder ohne Worte op. 19b, 1829-30).

1. INTRODUCTION

The word "explanation" means "to make plain", but derives from the Latin "to make into a *plane*", or "to flatten" [1]. In this paper, I will provide an information theoretic interpretation of this concept, where "flattening" takes on the meaning of reducing the weighting of less significant and accentuation of more significant data in a given analysis.

Zipf's law provides an excellent tool for this purpose, since it offers a measure of relative roughness of data sets [2]. The law itself states that data, which may come from many different contexts, tends to be approximately inversely proportional to the rank of that data. In other words, the second most common item occurs half as frequently as the most common, while the third occurs a third as often, and so forth.

Formally this can be stated as Eqn (1).

$$f(r) \sim \frac{d}{r^\alpha} \quad (1)$$

Where $f(r)$ denotes the frequency of an observation of rank r for a dataset with d as the largest observation. Zipf's law holds when the exponent α is approximately 1, a scenario that is also called pink noise. This can be contrasted with red and white noise, which has larger and smaller values for α respectively, corresponding to steeper and shallower slopes of the data when plotted against the rank on a log-log scale [3].

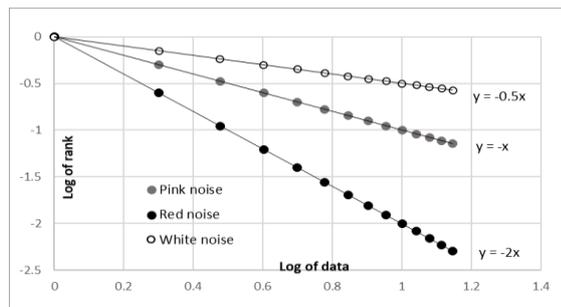


Figure 1. Comparison between three colors of noise on a log-log scale for data with a most common element of frequency 1.

Using this concept, the maximum amount of explanation offered by a data-set can now be stipulated as being a not quite vertical line, which corresponds to the most frequent element being maximally larger than all the following ones. The reason this still constitutes an explanation is that the slope of the remaining elements is relatively flattened compared to the most frequent elements in red noise data-sets, as can be seen in Figure 2.

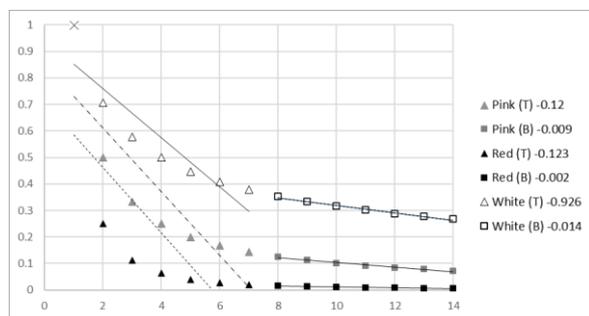


Figure 2. Linear comparison of noise profiles for pink ($\alpha = 1$), white ($\alpha = 0.5$) and red ($\alpha = 2$) data-sets. Trendlines are shown for the top (T) and bottom (B) subsets. Note that red noise shows the largest difference between the trendlines of the two subsets, with the steepest slope (-0.123) for the top portion and the flattest slope (-0.002) for the bottom portion. X denotes the common largest observation (1) for each sequence.

2. FROM RAW DATA TO ANALYSIS PRODUCT

2.1 Input Data

Noise profiles can be calculated for any dataset, which, in the case of music, can consist of notation, samples, listener evaluations or motion tracking data. For this paper, pitch class (PC) values and durations will be used, partly because they are exactly specified by composers, which makes them convenient as a neutral base to start an analysis from. But this should not be seen as precluding overlaying multiple layers of musical information as the techniques employed here are context-free. Indeed, this context-independence property is one of the main justifications for using power law representations as opposed to direct ones in music analysis.

For the present investigation, then, PCs are used as the basic data input, as opposed to absolute pitches. While this approach somewhat obscures structural analysis resulting from octavation, it allows a more precise focus on tonal domain information which will be the main subject of investigation here. Nevertheless, the techniques used here work in a comparable way regardless of the choice of raw data input.

2.2 Wavelet Analysis

2.1.1 Fourier Transforms

The mathematical antecedents of wavelet analyses can be traced back to the Fourier transform used by Joseph Fourier (1768-1830) to study heat equations [4]. The essential logic of the Fourier transform is to investigate a one-dimensional signal in terms of many overlapping sine wave functions, see Figure 3, in an analogous way to how Johann Fux's "Gradus ad Parnassum" [5] considers counterpoint at 1:1, 2:1 and 4:1 levels of analysis. Indeed, this overlapping level of analysis concept is found throughout music theory in general. The distinction between Fux and Fourier is that while Fux starts with a characteristic period of frequency overlap in mind, Fourier analysis starts with the data to discover the characteristic frequencies hidden within. This duality of frequency and data domains is a critical concept for wavelet analysis in general.

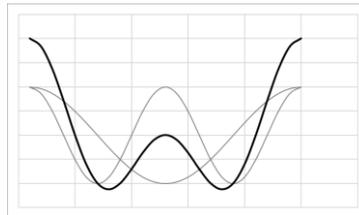


Figure 3. Sine waves can be added together to produce a compound waveform. The Fourier transform can be used to reconstruct the component sinusoids from a complex result.

2.1.2 Wavelets

The concept of analysis at multiple levels and across different domains can be further generalized to use waveforms other than the sine wave (called mother wavelets), including Haar wavelets, Daubechies wavelets or Coiflets, each with its own applications and features. The need for such alternative wavelets can be understood if one observes the difference between a piano note's waveform, which has a peak and decay, see Figure 4, with the sine wave model that rises and falls smoothly with each period. Unlike sine waves, wavelets have finite length and use stacked scaling analogously to the stacked frequencies in the Fourier transform. This allows for a better resolution of boundaries in a signal, for which the smooth, continuous curve of the sine wave is not ideal.

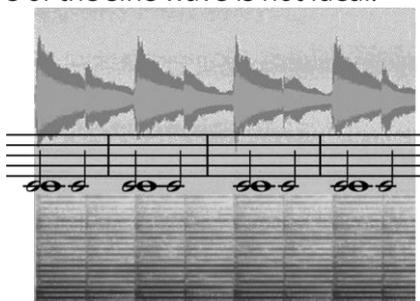


Figure 4. Amplitude waveform (top) music notation (mid) and spectrogram representations (bottom) of the same event. This shows how the spectrogram can represent the similarity between events while the waveform shows the difference. Moving between representations can give a more complete picture.

2.1.3 Timescapes

By the early 21st century this multiscale idea had returned back to the musical domain in the form of timescape or pitch scape analyses such as those of Gianpaolo Evangelista [6] Craig Sapp [7], Nicholas Cook [8], Joshua Neumann [9], and Robert Lieck and Martin Rohrmeier [10], among others. These kinds of analyses may assign numerical values for the fit of one variable (such as PCs) using assessment tools such as

Carol Krumhansl and Edward Kessler's key profiles [11] as mother wavelets. This allows key assessments to be made across multiple scales.

2.1.4 Power law wavelets

The approach in this paper is to use power law noise profiles, as defined by the α value of each scaled and translated time-slice, as the mother wavelet. This allows us to move back and forth between the data domain and the structure domain in the same way that Fourier analysis moves between time and frequency domains.

The main advantage of power laws in this context is in their scale invariance property, which is to say that multiplying all values in the data by some constant — i.e. changing the value of d in Eqn. (1) — does not change the noise profile. This not only considerably simplifies calculations but also allows for optimization algorithms to be performed on α values. Furthermore, because this approach is context independent, it simplifies cross-domain comparisons, for example between dynamics and inter-onset intervals, in addition to the PC approach taken here.

3. STRUCTURAL ANALYSIS – THREE SHORT WORKS FOR PIANO

3.1 “Von fremden Ländern und Menschen” – Robert Schumann

3.1.1 Rationale

While power law analyses of this type can be quite variable, they are also remarkably robust, preserving critical structural information in the face of quite extreme manipulation. This is an important feature that will be exploited later to “flatten” descriptions. However, in order to see power law wavelets in action, it is first necessary to perform a baseline structural analysis, so that there is something to compare the outputs to. These naïve analyses are placed above each wavelet analysis hereafter. Ideally the process should be an interactive one, so that the initial analysis informs how a power law output is interpreted, but also so that the output informs analysis.

3.1.2 Overview

The first of the three piano works under investigation, Robert Schumann's “Von fremden Ländern und Menschen” (no.1 from Kinderszenen op. 15, 1838), is instructive in this regard, because it is both structurally and tonally fairly transparent. The form can be described as rounded binary, AA || BABA, but it does not exhibit the usual modulation to the dominant expected in a traditional Minuet form. In fact, perhaps surprisingly given the title, the entire work stays in a G major region. The “strangeness” that justifies the characterization of “fremden” in the title is provided by an almost immediate secondary dominant (vii^{o7}) of V in the theme itself, which is repeated at the outset of each A section. This simple device very effectively offsets the tonal stability observed at the phrase level.

The only other chromatic alteration (which could be read as a chromatic passing note) is a raised dominant producing a III⁷ which leads to a bridge passage running up an e minor scale that ends on a D (the dominant of the G major home key), kicking off the main theme in a very slightly modified form again. Being two bars long, as opposed to the four bars of the main phrases, this bridge passage — denoted by * in Figure 5 — introduces an asymmetry into the second part which is resolved by the repetition of the section.

3.1.3 Raw Wavelet Analysis

Given the considerable amount of repetition employed by Schumann here, it is perhaps not surprising that the raw α value is quite large at 1.4 where Zipf's law would suggest an expected value of 1 (although it may be noted that frequent use of octave doublings will tend to redden PC noise profiles in all the works

in this paper). That means that even in raw, unadjusted form, the fractal nature of the work can be clearly observed in Figure 5 in the form of similar shapes at different scales (vertical displacements).

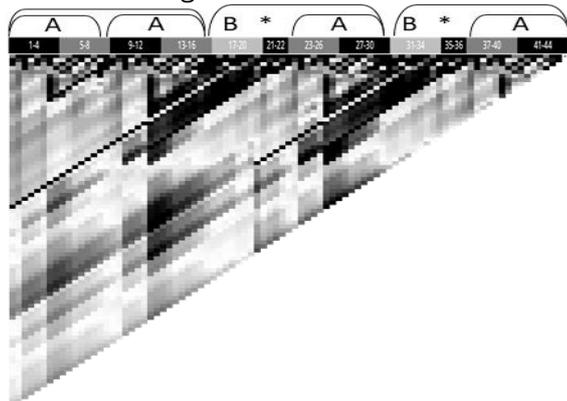


Figure 5 Raw power law wavelet analysis of Schumann's op. 15 no. 1. The analysis is arranged analogously to music, with shorter segments towards the top and longer ones to the bottom. This means that color changes on the vertical axis represent the beginning of sections and diagonal lines represent end-points for segments of different sizes. The color coding shows the 50th percentile α values as light and extreme high and low values (95th percentile) as dark.

3.1.4 The Resultant

The method adopted to adjust this α while moving the analysis into the tonal domain is to split the PC counts into twelve sets corresponding to the twelve keys. This results in twelve distinct α values for any segment of data, which can be recombined either as an average value or as a second order α of α value. These two combinations tend to imperfectly cancel out biases in each other, which makes the average of their values a useful measure in practice. I have called this value the "resultant".

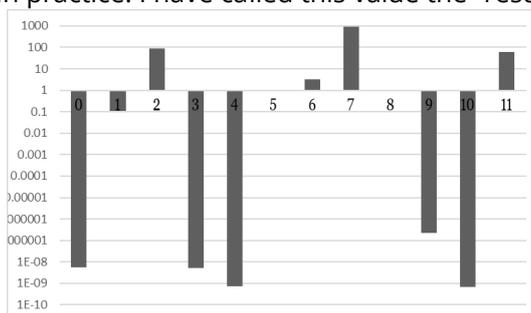


Figure 6. The weightings of each PC required to achieve a resultant α of 9. Note the logarithmic scale in the y-axis. PCs 5 and 8 do not occur in the work. Parallelisms can be observed between the leading note-tonic relations of both G major and D major, an important contrast in the main theme. A similar parallelism can be observed between PCs 3 & 4 and 9 & 10, reflecting a tenor voice chromatically descending three note motive found in the opening theme and the bridge passage respectively. PC 11 is very prominent in the main melody, while PC 0 plays a subsidiary, bridging role.

3.1.4 Resultant Wavelet Analysis

The resultant can be algorithmically optimized by biasing each note count with a calculated weighting. I have employed one of Microsoft Excel's built-in solver functions (GRG nonlinear) for this task, repeatedly

performing the operation in search of the highest whole number resultant α without any zero weights. This generates the weighting seen in Figure 6 which in turn generates the wavelet heatmap in Figure 7.

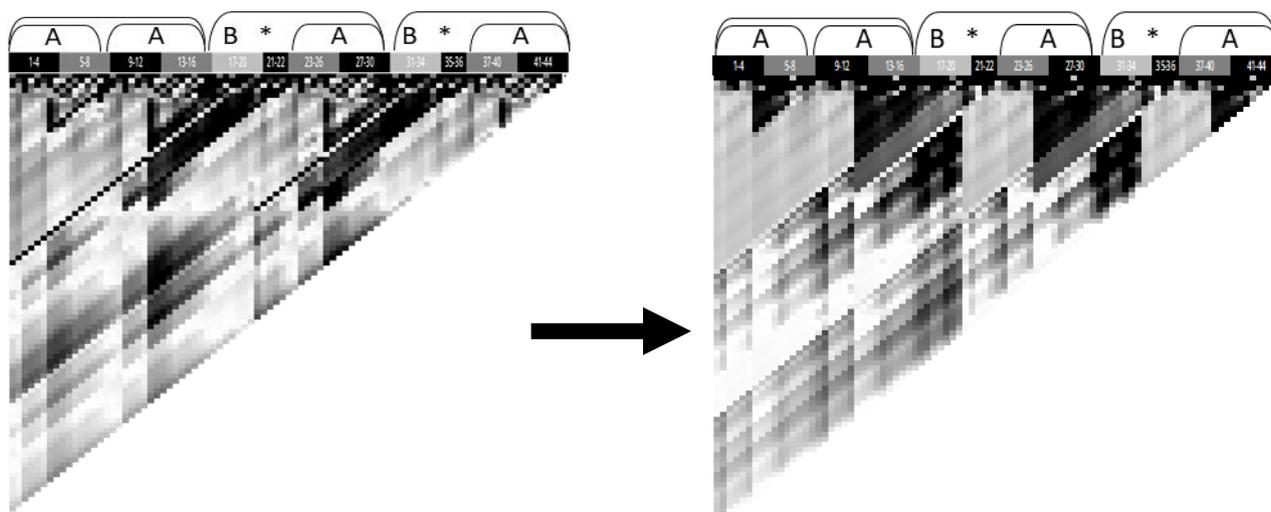


Figure 7. Power law analysis of Schumann's op. 15 no. 1 optimized for a resultant α of 9 (right). Observe how Schumann overlaps the subphrases, phrases and sentences by emphasizing the leading note tonic motion in G major and D major from the very first bar, resulting in a markedly self-similar fractal pattern.

3.2 Prelude no. 20 - Frédéric Chopin

3.2.1 Overview

Despite being one of the shortest of Chopin's 24 Preludes, in the words of James Huneker [12], the Prelude no. 20 "holds in only thirteen bars the suffering of a whole people, a whole nation."

In structural terms it is comprised of three sections ABB plus a final bar coda consisting of a single c minor chord. The first section is prominently subdivided into two parts by a marked transition between an f minor first inversion chord at the end of bar 2 to a G major chord at the start of bar 3.

The B sections are exact repetitions in pitch class terms, although they would be distinguished in performance by a quieter dynamic together with an *una corda* indication. Such distinction could be incorporated into a power law analysis by adjusting the weightings to follow dynamic and tempo changes, but for present purposes PCs are considered in isolation. As with the decision to use PC as opposed to absolute pitches for analysis, the omission of such details somewhat obscures some structural features and highlights the role of performers in bringing musical works to life beyond what a high-level analysis can reveal.

Regardless, the richness of the initial analytical output shows how much information can be contained in such a short and ostensibly simple work.

3.2.2 Raw wavelet analysis

For the Prelude, the raw α value is still larger than the expected Zipfian optimum of 1, but at a value of 1.39 it is a little closer than the 1.4 of the Schumann.

Major boundaries, such as the transition between the A and B sections do not appear to be as clearly represented in pitch class space as might be expected apart from very short period segments. However, there are clear segmentations visible at secondary locations in Figure 8, such as in the respective second bars of the B sections and in the region surrounding the dramatic harmonic shift between bars 2 and 3.

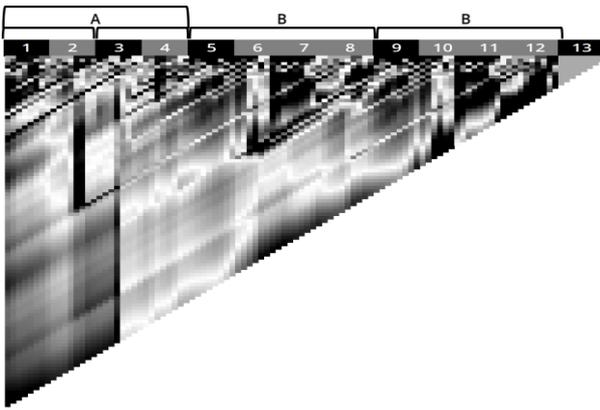


Figure 8. Raw power law wavelet analysis of Chopin's op. 28 no. 20.

3.2.3 Resultant

The solver optimization for the Chopin prelude arrived at a resultant of $\alpha = 7$ before pitches with non-zero count were assigned a zero weighting by the algorithm. Crucially, lower α values do not show markedly different patterns of weighting solution behavior even though, as α values go to extreme values away from optimum mean value of 1, the number of possible solutions is reduced, constraining the range of outcomes. This means that, while moving to a high α does not entirely eliminate the possibility of multiple solutions being available in the solution space, it does reduce the probability of finding alternatives, at least for a given total weighted PC count.

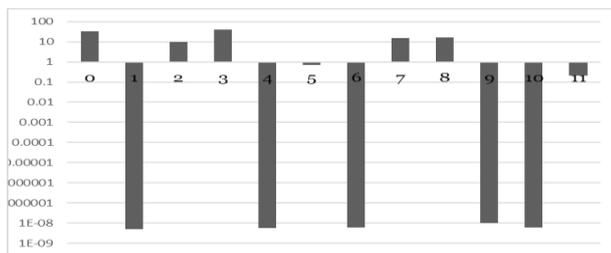


Figure 9. Weighting profile for Chopin's Prelude op. 28 no. 20. The final α obtained here was 7.

In the present case the solver solution shown in Figure 9 highlights the notes of c harmonic minor (the music-theoretically expected key) precisely. This suggests that, while the work features a lot of chromatic and dissonant harmony, the tonal center is quite robustly centered on the initial tonic. This reading is strongly supported by a consideration of the bass line, which moves chromatically but always within the c minor frame, even when considering the prominent use of the Neapolitan flattened degree at several points. This technique allows Chopin to evoke a strong sense of longing and nostalgia that is so characteristic of the work.

3.2.3 Resultant wavelet analysis

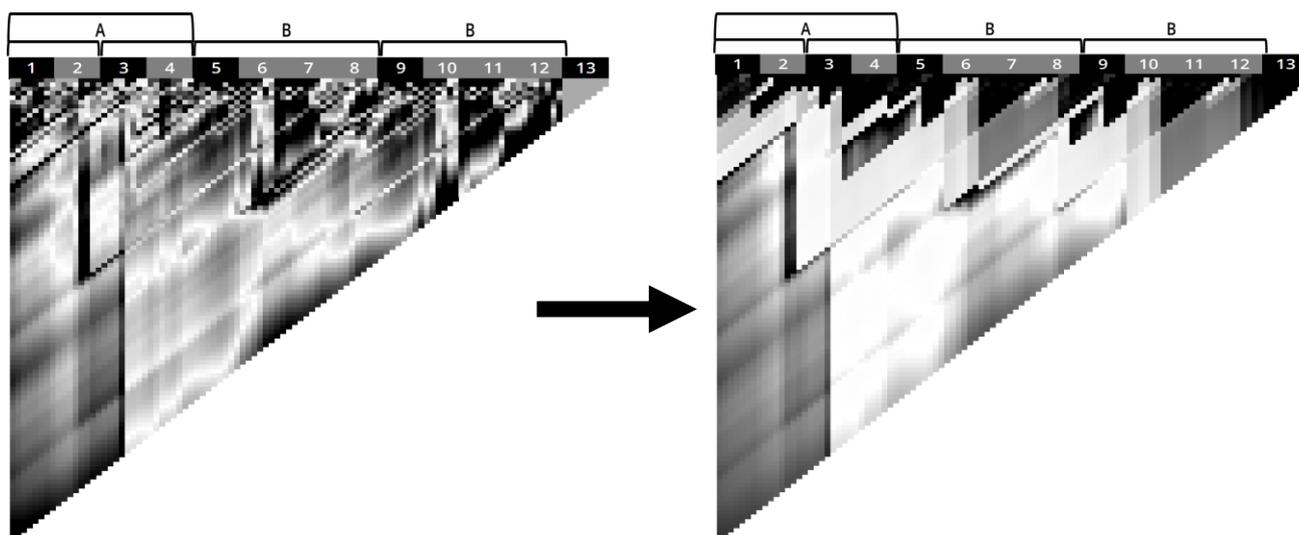


Figure 10. Power law analysis of Chopin's Prelude op. 28 no. 20 with a resultant $\alpha = 7$ weighting applied (right).

The wavelet analysis produced by the red noise resultant, Figure 10, displays a somewhat surprising amount of overlap between the sections. This is partly explained by the lack of sharp tonal distinctions between the sections at the precise boundaries. The change between section A and B, for example, despite being audibly distinct, involve a harmonically almost mundane V-I motion in c minor. The repetition of the B sections is even less harmonically marked, with the section beginning and ending on a c minor tonic chord. Perhaps counter-intuitively, Chopin uses this harmonic overlap to accentuate the abrupt dynamic and registral changes at these points (which are invisible to the PC analysis conducted here), perhaps as if to say "Even if I have gone to another place, I am still the same".

3.3 "Venetianisches Gondellied" - Felix Mendelssohn

3.3.1 Overview

Even though Mendelssohn's Venetian Boat Song op. 19b no. 6 work is the oldest of the three works in the current set, it is one of three "Songs Without Words" with this title, the other being op. 30 no. 6 from 1833-4 and op. 62 no. 5 1842-4, all of which have a similar style, texture and mood. As well as being the earliest, the op. 19b no. 6 Boat Song is also by far the most complex of the three works under investigation in this paper.

While this relative complexity is easy to accommodate by the method thus far described, it does highlight the fact that even minimally complex musical works can be remarkably information dense and analytically rewarding. This offers a salutary lesson in avoiding an over-strict reading in terms of semantically laden categories such as the tonal template employed here. Indeed, even though this work begins and ends in a fairly unambiguous g minor, the melody, which first appears in bar 3, has a distinctly d Phrygian mode characteristic which to some extent evades a strict twelve key-space interpretation.

This information density of music that so evades linguistic tractability led Mendelssohn to observe: "What the music I love expresses to me, is not thought too indefinite to put into words, but on the contrary, too *definite*" [13] (Mendelssohn's italics).

Seen in this way, any attempt to seek a red-noise solution algorithmically is, even if descriptively useful, a somewhat destructive act, since it cleans up and reduces details that are artistically indispensable.

3.3.2 Raw wavelet analysis

The raw analysis of the Boat Song reveals an α value of 1.589, which is somewhat surprising, in that a value closer to the Zipfian optimum value of 1 was expected. The higher value perhaps reflects the fact that the

complexity here is mostly textural, and indeed structural in nature, which would not present itself to a global tonal analysis.

This textural complexity also highlights the difficulty in supplying works of this style with a music-theoretically informed one-dimensional structural analysis. The problem is essentially one of having to decide whether structural boundaries should be placed at the beginning of preceding bridging or introductory material or where thematic material associated with a section is actually introduced. Part of the problem is that it is often considered good compositional practice to obscure such obvious boundaries, especially in the Romantic style of the 19th century.

With this in mind I have opted to faithfully reproduce my naïve first pass score based structural analysis as reference, see Figure 11.

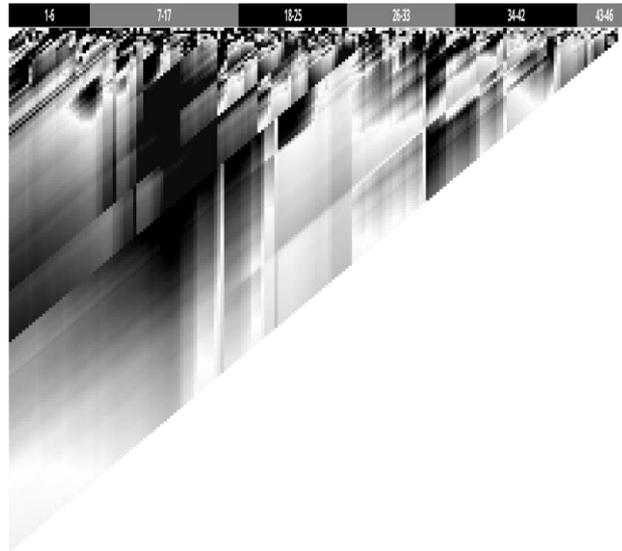


Figure 11. Raw wavelet analysis of Mendelssohn's "Venetianisches Gondellied" (no. 6 from Lieder ohne Worte op. 19b).

A high-level overview reveals that, for the raw analysis, while there are sharp boundaries, they tend to not correspond exactly to thematic entries, but tend to correspond rather better to bridging, and even cadential material. The correspondence of this analytical perspective to naïve listener experience is an interesting subject for further study, as is its generalizability to works by other composers and in other genres.

3.3.2 Resultant

Solver solutions represent local minima, but cannot discover global minima. As such, it is never possible to entirely eliminate the possibility that different starting conditions would result in completely different weightings which, to the extent that the weightings printed throughout this paper are musically interpretable, may not make any equivalently discernable music-theoretic sense.

To that end, throughout the solutions offered here, the starting point for the algorithm has consistently been a flat weighting (i.e., with all values set to 1). Additionally, a sanity check is provided by the simple counts of pitches in the works, which typically provide approximate analogues to the solver derived weighting profiles.

With that in mind, and since the Mendelssohn Boat Song already starts out with a high α , a rather more extreme solution was sought in this case. This was achieved by incrementally increasing the α by the same method as previously until a value of 8.015 was reached (again, representing the limit before zero weightings were obtained).

A higher α can, however, always be forced by replacing any of the individual weights by some large number. To account for this, each weight was successively replaced by a value of 10^5 , an operation that yielded a new set of weightings with their own profile which was similar, but larger than the original. Since all the α values were above 8, each weight had 8 subtracted from in and the resulting profile was used as

a new initial seed. This time the multi-start option was enabled and the process of iteratively driving the solver to higher levels of α was re-initiated, this time reaching a final value of 30 before zeros appeared in the weighting solution. The result is shown in Figure 12.

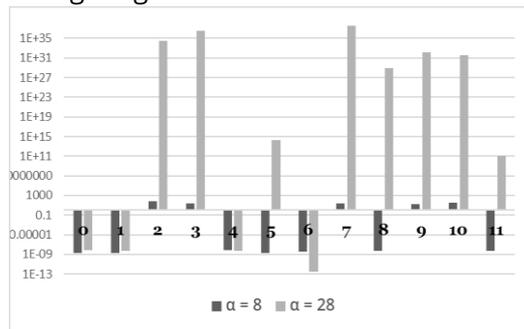


Figure 12. Comparison between the initial weighting result and a second order result using a more intensive multi-start process with a start point derived from the original weighting.

The result produced by this process is not clearly superior to the less intensive approach that yielded the initial α of 8. In the latter case, PCs {2, 3, 7, 9, 10} were emphasized, corresponding quite well to a partial g minor or d Phrygian mode, only missing out the expected PCs 0 and 5 which would have completely described g natural minor. The extreme α solution, meanwhile, exaggerated the same PCs, but with the addition of the somewhat unexpected PCs 8 and 11, in addition to the initially downranked PC 5. A closer inspection of the score reveals that PCs 8 and 11 only appear near the middle of the work at bars 18 to 20, making them important, albeit not globally prominent, structural markers in their own right. PC 5, meanwhile only features prominently in the opening and closing theme, perhaps explaining why these moments were somewhat better differentiated in the $\alpha = 28$ analysis.

3.3.3 Resultant wavelet

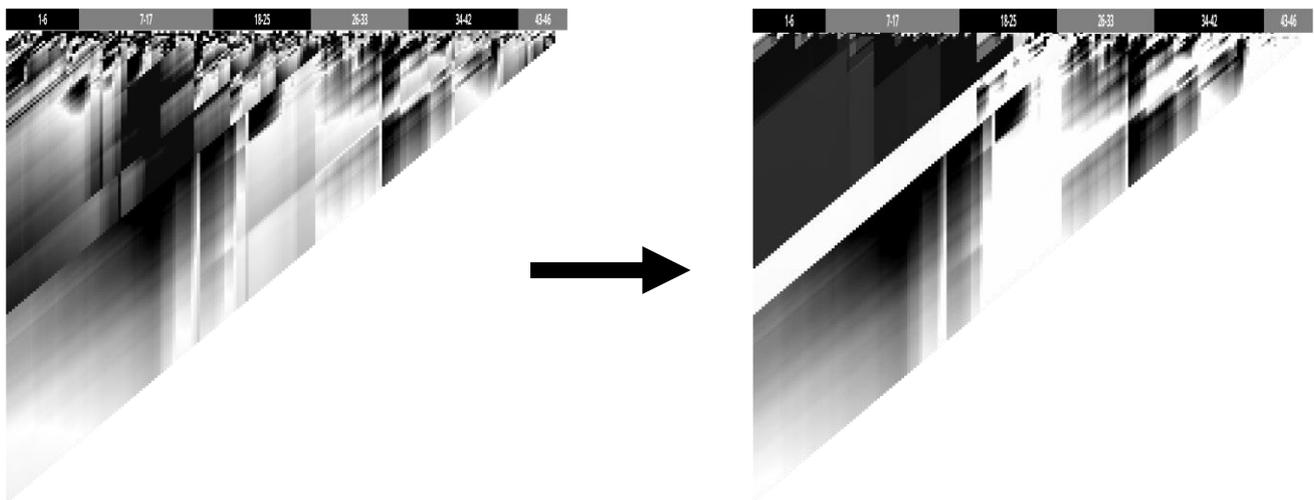


Figure 13. Weighted power law analysis of Mendelssohn's "Venetianisches Gondellied" (right). The differences between the $\alpha = 8$ and $\alpha = 28$ weighting are important, but small. The $\alpha = 28$ profile is shown here.

Despite the large apparent differences in the weighting profiles, the higher α weighting seen in Figure 13 offers only very marginal improvement over the lower value. This suggests that there is some tradeoff for using a secondary starting point, with slight improvements in grain being offset by a somewhat less informative weighting profile. Future work may involve testing more extreme adverse starting points to fully explore the robustness of the technique.

4. CONCLUSION

The main conclusion that can be drawn from this overview is that, with caveats, driving a resultant to red noise by algorithmically adjusting weighting is generally robust, reliable and informative. The results, however, cannot be naively interpreted, because music is potentially unlimited in its semiotic dimensionality. It is therefore always necessary to refer back and forth between musical, physical, cultural or semantic knowledge, and analysis products. This follows from the fact that, if there are significant features that are not in the scope of the initial counting protocol, no amount of classical information theoretic analysis will be able to recover them. This follows naturally from the principle that reductive counting techniques (such as PC encoding) destroys information. Intriguingly, though, this may not be an absolute prohibition, since the no-hiding theorem, which was experimentally confirmed in 2011 [14], says that information cannot be destroyed in the quantum domain. This offers an intriguing application for qubit (quantum bit) representation with regards to computational music analysis which may be explored in future studies.

5. REFERENCES

- [1] D. Harper. "Explanation." Online etymology dictionary. <https://www.etymonline.com/word/explanation>. Accessed 15 May 2021.
- [2] D. Kim, S. Son, and H. Jeong, "Large-Scale Quantitative Analysis of Painting Arts". *Scientific Reports*, vol. 4, no. 7370, 2014.
- [3] D. L. Gilman, F. J. Fuglister, and J. M. Mitchell Jr., "On the Power Spectrum of 'Red Noise'", *Journal of the Atmospheric Sciences*, vol. 20, no. 2, 1963.
- [4] J. Fourier, *Théorie Analytique de la Chaleur*. Paris, France: F. Didot, 1822.
- [5] J. Fux, *Gradus ad Parnassum*. Vienna: Johann Peter van Ghelen, 1725.
- [6] G. Evangelista, "Pitch-Synchronous Wavelet Representations of Speech and Music Signals", *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3313-3330, Dec. 1993.
- [7] C. Sapp, "Visual Hierarchical Key Analysis", *ACM Computers in Entertainment*, vol. 4, no. 4, 2005.
- [8] N. Cook, "Performance Analysis and Chopin's Mazurkas", *Musicae Scientiae*, vol. 11, no. 2, pp. 183-207, 2007.
- [9] J. Neumann, "Phenomena, Poiēsis, and Performance Profiling: Temporal-Textual Emphasis and Creative Process Analysis in Turandot at the Metropolitan Opera." *Empirical Musicology Review*, vol. 12, no. 3, pp. 247.
- [10] R. Lieck and M. Rohrmeier, "Modelling Hierarchical Key Structure with Pitch Scapes", in *Proc. of the 21st Int. Society for Music Information Retrieval Conf.*, Montréal, Canada, 2020, pp. 811-818.
- [11] C. L. Krumhansl and E. J. Kessler, "Tracing the Dynamic Changes in Perceived Tonal Organization in a Spatial Representation of Musical Keys." *Psychological Review*, vol. 89, no. 4, 1982, pp. 334-368.
- [12] A. Cassela, Ed. *F Chopin Preludi*, Milano, IT: Edizioni Curci, 1947.
- [13] F. Mendelssohn, *To Marc Andre Soucha, Lubeck*, in *Selected letters of Mendelssohn*, W. F. Alexander, Ed. New York: Macmillan and Co., 1894, p. 124.
- [14] J. R. Samal, A. K. Pati, and A. Kumar, "Experimental Test of the Quantum No-Hiding Theorem", *Physical Review Letters*, vol. 106, no. 8, Feb. 2011.