

Zipf's law and the creation of musical context

Damián H. Zanette

Consejo Nacional de Investigaciones Científicas y Técnicas
Instituto Balseiro
8400 San Carlos de Bariloche, Río Negro, Argentina
zanette@cab.cnea.gov.ar

Abstract

This article discusses the extension of the notion of *context* from linguistics to the domain of music. In language, the statistical regularity known as Zipf's law - which concerns the frequency of usage of different words - has been quantitatively related to the process of text generation. The connection is established by Simon's model, on the basis of a few assumptions regarding the accompanying creation of context. This model captures the essential mechanism of repetition of perceptual elements, which underlies the construction of a structured, comprehensible message. Here, it is shown that the statistics of note usage in musical compositions are compatible with the predictions of Simon's model. This result, which gives objective support to the conceptual likeness of context in language and music, is obtained through automatic analysis of the digital versions of several compositions. As a by-product, a quantitative measure of context definiteness is introduced, and used to compare tonal and atonal works.

Resumen (Spanish)

En este artículo se discute la extensión de la noción lingüística de *contexto* al dominio de la música. En el lenguaje, la regularidad estadística conocida como ley de Zipf - referida a la frecuencia de uso de diferentes palabras - ha sido vinculada cuantitativamente con el proceso de generación de un texto. La relación queda establecida por el modelo de Simon, en base a unas pocas suposiciones sobre la simultánea creación de contexto. Este modelo capta el mecanismo esencial de repetición de elementos perceptivos, que subyace en la construcción de un mensaje

estructurado e inteligible. Aquí se muestra que la estadística del uso de notas en composiciones musicales es compatible con las predicciones del modelo de Simon. Este resultado, que afirma objetivamente la afinidad conceptual del contexto en lenguaje y en música, se obtiene mediante el análisis automático de la versión digital de varias composiciones. Como subproducto, se introduce una medida cuantitativa del grado de definición del contexto, la cual es usada para comparar obras tonales y atonales.

1. Introduction

The appealing affinity between the cognitive processes associated with music and language has always motivated considerable interest in comparative research (Patel, 2003). Both music and language are highly structured human universals related to communication, whose acquisition, generation, and perception are believed to share at least some basic neural mechanisms (Maess *et al.*, 2001). The analysis of these concurrent aspects has naturally led to the attempt of extending concepts and methods of linguistics to the domain of musical expression. Linguistically-inspired approaches to music, however, have not always taken into account the crucial difference of nature between the information conveyed by music and language. Consequently, they have often remained at the level of a metaphoric parallelism (Bernstein, 1973). A scientifically valuable comparative investigation of music and language should begin by an accurate definition of common concepts in both domains. In this direction, systematic theoretical and empirical exploration of the concepts of musical grammar (Lerdahl and Jackendoff, 1983), syntax (Patel, 2003), and semantics (Koelsh *et al.*, 2004) has been undertaken.

In this article, I explore the possibility of extending to the domain of music a quantitative feature of language, related to the frequency of word usage - namely, Zipf's law. The significance of Zipf's law for language turned out to be a controversial matter in the past (Simon, 1955; Mandelbrot, 1959). However, the most successful explanation of Zipf's law - given by Simon's model - is based on linguistically sensible assumptions, associated with the mechanisms of text generation

and the concept of context creation (Simon, 1955; Montemurro and Zanette, 2002; Zanette and Montemurro, 2004). This supports the assertion that Zipf's law is relevant to language. Moreover, since it involves a quantitative property, an extension to the domain of music can, in principle, be precisely defined.

Zipf's law has already been studied in music from a phenomenological perspective, without reference to any possible connection between linguistics and music theory (Boroda and Polikarpov, 1988; Manaris *et al.*, 2002, 2003). The main aim of this article is to discuss Zipf's law as a by-product of the creation of musical context, attesting the validity of extending the assumptions of Simon's model to music. I begin by reviewing the formulation of Zipf's law and Simon's model for language, with emphasis in their connection with the concept of context. Then, I discuss the extension of this concept to music. Finally, I show with quantitative examples that Simon's model can be successfully applied to musical compositions, which provides evidence of analogous underlying mechanisms in the creation of context in language and music. Context thus arises as a well-defined common concept in the two domains.

2. Zipf's law and Simon's model in language

In the early 1930s, G. K. Zipf pointed out a statistical feature of large language corpora - both written texts and speech streams - which, remarkably, is observed in many languages, and for different authors and styles (Zipf, 1935). He noticed that the number of words $w(\underline{n})$ which occur exactly \underline{n} times in a language corpus varies with \underline{n} as $w(\underline{n}) \sim 1/\underline{n}^\gamma$, where the exponent γ is close to 2. This rule establishes that the number of words with exactly \underline{n} occurrences decreases approximately as the inverse

square of \underline{n} . Zipf's law can also be formulated as follows. Suppose that the words in the corpus are ranked according to their number of occurrences, with rank $\underline{r} = 1$ corresponding to the most frequent word, rank $\underline{r} = 2$ to the second most frequent word, and so on. Then, for large ranks, the number of occurrences $\underline{n}(\underline{r})$ of the word of rank \underline{r} is given by $\underline{n}(\underline{r}) \sim 1/\underline{r}^{\underline{z}}$, with \underline{z} close to 1. The number of occurrences of a word, therefore, is inversely proportional to its rank. For instance, the 100-th most frequent word is expected to occur roughly 10 times more frequently than the 1000-th most frequent word. Figure 1 illustrates Zipf's law for Charles Dickens's David Copperfield. All its different words have been ranked, the number of occurrences \underline{n} of each word has been determined, and \underline{n} has been plotted against the rank \underline{r} . In this double-logarithmic plot, straight lines correspond to the power-law dependence between \underline{n} and \underline{r} reported by Zipf.

FIGURE 1 NEAR HERE

Zipf himself advanced a qualitative explanation for the relation between word frequency and rank, based on the balanced compromise between the efforts invested by the sender and the receiver in a communication process (Zipf, 1949). A quantitative derivation of Zipf's law was later provided by H. A. Simon, in the form of a model for text generation (Simon, 1955). The basic assumption underlying Simon's model is that, as words are successively added to the text, a context is created. As the context emerges, it favours the later appearance of certain words - in particular, those that have already appeared - and inhibits the use of others. The model aims at capturing the essentials of the mechanism which, by repetition of linguistic elements - in this case, words - is at work in the construction of a structured,

comprehensible text. The repetition of perceptual elements is one of the basic ingredients in the conception of intelligible structures and in the ensuing cognitive response to their reception, including the creation and evocation of memories (Brown and Hagoort, 2000). In section 3, this mechanism is invoked again in connection with the creation of musical patterns.

In its simplest form, Simon's model postulates that, during the process of text generation, those words that have not yet been used are added at a constant rate, while a word that has already appeared is used again with a frequency proportional to the number of its previous occurrences. These simple rules are enough to prove that, in a sufficiently long text, the number $w(n)$ of words with exactly n occurrences is, as noticed by Zipf, $w(n) \sim 1/n^\gamma$. The exponent γ is determined by the rate at which new words are added. It takes the value $\gamma \approx 2$, as observed in real texts, when that rate is close to zero. Thus, according to Simon's model, Zipf's law is realised when the generation of a text reaches stages where the frequency of appearance of new words has declined sufficiently.

Simon's model can be refined by assuming that, as it is observed to happen in real texts, the rate of appearance of new words varies as the text becomes longer (Montemurro and Zanette, 2002; Zanette and Montemurro, 2004). A phenomenological description of this variation is specified by relating the number V of different words, *i.e.* the lexicon size, to the length T of the text, as $V \sim T^\nu$. The exponent ν can be identified with the ratio between a relative variation $\Delta V/V$ in the lexicon size to the corresponding relative variation $\Delta T/T$ in the text length. A small value for ν corresponds to a lexicon whose size increases slowly as compared with the

text growth, and *vice versa*. Introducing the above relation between \underline{V} and \underline{T} in Simon's model, the number of words with exactly \underline{n} occurrences turns out to be $\underline{w}(\underline{n}) \sim 1/\underline{n}^{1+\nu}$.

A further refinement of Simon's model consists in assuming that there is an upper limit \underline{n}_0 for the number of occurrences of any single word. This reasonable restriction implies that no word can appear more than \underline{n}_0 times in the text. Consequently, the function $\underline{w}(\underline{n})$ drops abruptly to zero for $\underline{n} > \underline{n}_0$. With these two new ingredients, Simon's model predicts that the number of occurrences as a function of the rank is

$$\underline{n}(r) = \frac{1}{(a+br)^{\underline{z}}} \quad (1)$$

with $\underline{z} = 1/\nu$. The constants \underline{a} and \underline{b} are given in terms of the parameters \underline{n}_0 and \underline{V} as $\underline{a} = 1/\underline{n}_0^\nu$ and $\underline{b} = (1-1/\underline{n}_0^\nu)/\underline{V}$. The upper limit \underline{n}_0 is in turn connected to \underline{V} and \underline{T} through the relation $\underline{T}/\underline{V} = \nu(\underline{n}_0^{1-\nu} - 1)/(1-\nu)(1-1/\underline{n}_0^\nu)$. For sufficiently large ranks, the form of $\underline{n}(r)$ given in equation (1) reproduces the expected "Zipfian" behaviour $\underline{n}(r) \sim 1/r^{\underline{z}}$. For small ranks, on the other hand, $\underline{n}(r)$ approaches a constant value and, consequently, the curve of \underline{n} versus r abandons the "Zipfian" regime and exhibits a shoulder-like profile. The dotted curve in figure 1 is a least-square fitting of the data of \underline{n} vs r for David Copperfield with equation (1). The remarkable agreement between the data and the fitting supports the hypotheses of Simon's model.

Thus, Simon's model interprets Zipf's law as a statistical property of word usage during the creation of context, as a text is progressively generated. As discussed in section 3, the creation of context in language is intimately related to the semantic contents of words, *i.e.* to their meaning. Semantics is in fact essential to the function

of language as a communication system. To assess the significance of Zipf's law and of the assumptions of Simon's model in the framework of music one must first examine to which extent the concepts of context and semantic contents can be extended to musical expression.

3. Semantics and context in music

In contrast to language, music lacks functional semantics.¹ Generally, the musical message does not convey information about the extra-musical world and, therefore, a conventional correspondence between musical elements and non-musical objects or concepts (*i.e.*, a dictionary) is irrelevant to its cognitive function. Unless music is accompanied by a text and/or by theatrical action, its semantic contents is usually limited to the onomatopoeic-like episodes of "musical pictures" or to a rather rough outline of mood, frequently determined just by rhythm and tonality. Assigning extra-musical meaning to a musical message is basically an idiosyncratic matter, yielding highly non-universal results.

On the other hand, the notion of context is essential to both language and music. In the two cases, context can be defined as the global property of a structured message that sustains its coherence or, in other words, its intelligibility (van Eemeren, 2001). Thus, such notion lies at the basis of the cognitive processes associated with written and spoken communication and with musical expression and perception. A long chain of words - even if they constitute a grammatically correct text - or a succession of musical events - even if they form, for instance, a technically acceptable harmonic progression - would result incomprehensible if they do not succeed at defining a

contextual framework. It is in this framework, created by the message itself, that its perceptual elements become integrated into a meaningful coherent structure.

In language, context emerges from the mutually interacting meanings of words. As new words are successively added to a text or speech stream, context is built up by the repeated appearance of certain words or word combinations, by the emphasis on some classes of nouns and adjectives, by the choice of tense, etc. These elements progressively establish the situational framework defined by the message in all its details. Thus, linguistic context is a collective expression of the semantic contents of the message, arising from the multiple structured relations between language elements.

In music, context is determined by a hierarchy of intermingled patterns occurring at different time scales. For the occasional listener, the most evident contribution to musical context originates at the level of the melodic material, whose repetitions, variations, and transpositions shape the thematic base of a composition (Schoenberg, 1967). The tonal and rhythmic structure of melody motifs and phrases constitutes the substance of musical context at that level. At larger scales, the recurrence of long sections and certain standard harmonic progressions determine the musical form. Crossed references between different movements or numbers of a given work establish patterns over even longer times. Meanwhile, at the opposite end of time scales, a few notes are enough to determine tempo, rhythmic background, and tonality (Krumhansl, 1990), through their duration and pitch relations.

An obvious difficulty in modelling the creation of musical context along the lines discussed in section 2 for language, which are based on the statistics of word usage, resides in the fact that the notion of word cannot be unambiguously extended to music (Boroda and Polikarkov, 1988). In language, words - or short combinations of words - stand for the units of semantic contents, with (almost) unequivocal correspondence with objects and concepts. Moreover, in the symbolic representation of language as a chain of characters, *i.e.* as a written text, words are separated by blank spaces and punctuation marks, which facilitates their identification - in particular, by automatic means. Music, on the other hand, does not possess any conventionally defined units of meaning. The notion of word is however conceivable in music by comparison with the linguistic role of words as “units of context”, namely, as the perceptual elements whose collective function yields coherence and comprehensibility to a message. In music, the role of “units of context” is played by the building blocks of the patterns which, at different time scales, make the musical message intelligible. Yet, the identification of such units in a specific work may constitute a controversial task. In previous studies on Zipf’s law in music, many musical elements - ranging from single notes to melodic digrams and trigrams, or doublets and triplets of pitch intervals between contiguous notes - have been proposed and employed as “units of context” (Manaris *et al.*, 2002, 2003).

In the quantitative investigation of context creation in music, I have chosen as “units of context” the building blocks of the smallest-scale patterns, namely, single notes. A note is here characterised by its pitch (*i.e.*, its position on the clef-endowed staff) and type (*i.e.*, its duration relative to the tempo mark). Its volume, timbre, and actual frequency and duration are disregarded. The contribution of notes to the creation of

musical context, determining tonality and the basis for rhythm, is particularly transparent (Krumhansl, 1990). In addition, the choice of single notes has several operational advantages. In the first place, the collection of notes available to all musical compositions - or, at least, to all those compositions that can be written on a staff using the standard note types - is the same. This collection of notes plays the role of the lexicon out of which the message is generated. Secondly, single notes are well-defined entities in any symbolic representation of music, either printed on a staff or in standardised digital formats, such as the Musical Instrument Digital Interface (MIDI). This makes possible their automatic identification, which, as described later, constitutes a crucial step in the analysis. Moreover, in order to extract any meaningful information from a statistical approach, it is necessary to work with relatively large corpora. The compositions used in the present investigation contain, typically, several thousand single notes. This figure remains well below the number of words in any literary corpus, which usually reaches a few hundred thousands (*cf.* figure 1), but is already suited for statistical manipulations.

The convenience of choosing single notes as the “units of context” is best appraised by comparing with other possible choices. Consider, for instance, a definition of “unit of context” in terms of pitch intervals between two successive notes (Manaris *et al.*, 2002). With this choice, the analysis becomes artificially restricted to those compositions where different voices are distinctly identifiable, so that the notion of *successive notes* is well defined. Alternatively, consider a definition in terms of melodic motifs. First of all, the limits of a melodic motif cannot be unambiguously determined. Furthermore, unless one takes into account the infinitely vast universe of all possible melodies, melodic motifs do not constitute a common lexicon for different

compositions. Finally, since melodic motifs are subject to modulation and variation as a work progresses, their automatic identification would demand resorting to the sophisticated computational procedures.

4. Application of Simon's model to music

The starting point in the study of the relevance of Simon's model to the creation of musical context, is Zipf's analysis of note usage. For a given musical composition, I have extracted information on note pitch and duration from a MIDI version of the work, following the procedure explained in the appendix.² As a result of this procedure, a list of different notes, characterised by their pitch and duration, and ranked according to their number of occurrences, was obtained. I denote by T the total number of notes (*i.e.*, the "text length", *cf.* section 2) and by V the number of different notes (*i.e.*, the "lexicon size").

I have performed Zipf's analysis on a variety of western music works, from different periods, styles, and with different musical forms. In this article, I present results for four compositions for keyboard, which insures a certain degree of idiomatic homogeneity in spite of the diversity of style. They are the Prelude N. 6 in d from the second book of *Das Wohltemperierte Klavier*, by J. S. Bach; the first movement, *Allegro*, from the Sonata in C (K. 545) by W. A. Mozart; the second movement, *Menuet*, from the Suite Bergamasque by C. Debussy; and the first of Three Piano Pieces (Op. 11, N. 1) by A. Schoenberg (for score sources, see the appendix). In all cases, I have disregarded short grace notes, which have not been written down by the composer and whose realisation relies on the performer, and have not taken into

account full-section repetitions, which contribute to musical context at the largest time scales only.

FIGURE 2 NEAR HERE

Figure 2 shows, as full dots, the number of occurrences \underline{n} versus the rank \underline{r} for single notes in the four works listed above. The respective values of \underline{V} and \underline{T} are indicated in each panel. The merest inspection of these Zipf's plots reveals a striking similarity in the functional shape of $\underline{n}(\underline{r})$ for the four data sets. I have obtained the same kind of shape for all the compositions analysed following Zipf's prescription. This similarity already suggests the existence of a common underlying mechanism, determining the relative frequency at which different notes are used, independent of work length, musical form, tonality, style, and author.

Note that, in contrast to figure 1, the plots of figure 2 lack the long linear regime corresponding to the power-law dependence of $\underline{n}(\underline{r})$. This circumstance, which can be ascribed to the relatively minute values of \underline{V} and \underline{T} for musical compositions as compared with literary corpora, does not preclude the application of Simon's model. In fact, according to equation (1), the "Zipfian" regime is attained for sufficiently large ranks only. The empirical data obtained from Zipf's analysis of note usage must be rather compared with the full form of $\underline{n}(\underline{r})$, as given by equation (1).

Curves in figure 2 stand for least-square fittings of the data with equation (1). As discussed in section 2, the constants \underline{a} and \underline{b} can be calculated in terms of the respective values of \underline{V} and \underline{T} , and of the exponent ν . Consequently, ν is the only free

parameter to be determined by the fitting. The resulting values of the exponent ν are quoted in figure 2. The agreement between the empirical data and the prediction of Simon's model is remarkably good for the four data sets. A chi-square test of the quality of fitness validates the hypothesis that these data are statistically equivalent to equation (1) at a confidence level close to 100 %. This implies that the results of Zipf's analysis are compatible with the hypothesis that single-note usage follows the assumptions of Simon's model. Specifically, they are in agreement with the assumption that the occurrences of a given note promote its later appearance, with a frequency that grows as the number of previous occurrences increases. According to the above discussion, this process stands for the basic mechanism of context formation.

While the four data sets shown in figure 2 are consistent with Simon's model and, in fact, display a common functional dependence between \underline{n} and \underline{r} , a quantitative disparity becomes apparent by comparing the respective values of the exponent ν , obtained from the least-square fitting. Recall from section 2 that this exponent quantifies the functional relation between the lexicon size \underline{V} and the text length \underline{T} , as $\underline{V} \sim \underline{T}^\nu$. Small exponents are an indication of a compact lexicon, determining a robust context that remains relatively stable and well defined as the text progresses. On the other hand, large exponents reveal an abundant lexicon, related to a ductile, unsteady, more tenuously defined context. In terms of context, therefore, the exponent ν can be interpreted as quantitative measure of variability or, conversely, of definiteness.

In the four musical works analysed here, the exponent ν happens to grow chronologically, following the composition dates. Its variation from Bach to Debussy

is however not significant. In fact, the analysis of other keyboard works by Bach and Mozart - for instance, other preludes from *Das Wohltemperierte Klavier* and other sonatas - yields values between 0.25 and 0.45. The only significant difference corresponds therefore to Schoenberg's Piano Piece. This work is well known as a landmark of consistent atonality, where the construction of a tonal context has been avoided on purpose (Perle, 1991). The absence of one of the contextual elements determined at the level where single notes act as "units of context" is clearly manifested by the large value of v resulting from the present analysis.

5. Conclusion

The extension of the notion of semantic contents from language to music is jeopardised by the fact that musical elements do not possess functional meaning in the extra-musical world. Semantics, in contrast, is essential to the function of language. On the other hand, context - whose role in language is closely related to semantics - stands for a significant feature common to linguistic and musical messages. In both domains, context denotes a property emerging from the interaction of the perceptual elements that compose the message, that makes the message intelligible as a whole. The nature of the information borne by music differs substantially from that of language. However, the combination of those elements in a hierarchically organised sequence, whose structure sustains its comprehensibility, lies at the basis of the creation of context in the two domains.

In this article, I have provided evidence supporting the assertion that the definition of linguistic context can be shared with music. Fortunately enough, context can be

conceptually related to a quantitative property of literary corpora, enunciated by Zipf's law, whose validity in a musical corpus can be investigated by objective means. It is Simon's model which establishes the connection between message generation, context creation, and Zipf's law. For language, Simon's model postulates that the rate at which a given word is used in a text is proportional to the number of its previous appearances. It quantifies the mechanism of repetition of perceptual elements in the creation of intelligible patterns, whose structure elicits the cognitive functions that lead to their comprehension. In music, this mechanism is inherent to the more basic principles of exposition of compositional material, such as motive and melodic construction, rhythm and tonality definition, and harmonic progression.

The evidence arises from the confirmation that musical corpora verify the predictions of Simon's model, an approach that relies on purely mathematical operations. As a by-product, this approach yields a quantitative measure of context definiteness - the exponent v . A demonstration of this measure has been drawn from the comparison of an atonal musical work with tonal compositions: in the former, the absence of tonal context results in a much larger value of v .

Of course, the present mathematical approach is not assumption-free. In particular, a crucial choice was made at the moment of extending the notion of word to musical messages. It would be interesting to consider alternative extensions, at the level of melodic motifs, harmonic sequences, or rhythmic patterns, and thus explore the concept of musical context at different scales.

Appendix. Music sources and methodological aspects

The scores of the musical compositions studied in this paper were obtained from the following sources: J. S. Bach, *The Well-Tempered Clavier (Books I and II, Complete)*, Dover, New York, 1983; W. A. Mozart, *Sonate e Fantasia*, Ricordi, Milan, 1949; C. Debussy, *Suite Bergamasque (II Menuet)*, www.music-scores.com; A. Schoenberg, *Three Piano Pieces (Op. 11)*, www.colleges.org/~music/modules/op11. The two latter sources provide free-share music scores.

MIDI files were created with a commercial music editor, Mozart Virtuoso Edition, version 3.0.4.0 (www.mozart.co.uk), and processed with a free-share MIDI assembler/disassembler, `mididsm.exe` (www.borg.com/~jglatt/). This software produces, from a MIDI source file, a text file with a list of MIDI events in the source. Each note-on or note-off event includes, among other data, information on the time at which the event occurs, and on the pitch of the note (Lehrman and Tully, 1993). This information was automatically read from the text file - discarding data on volume, MIDI channel, etc. - and note durations were calculated by subtracting the times corresponding to successive note-on and note-off events.

The output of the above procedure was a list of notes (or “words”), each of them given by a pitch-duration pair. Before processing this list, however, it was necessary to overcome a problem in the definition of note duration. Presumably due to quality requirements in the production of the sound sequence, the music editor used to create the MIDI files assigns to each sound a duration which is slightly shorter than the corresponding note type. For instance, in common (4/4) time, a crotchet lasts 89 of the

96 time units in which each beat is divided by the MIDI editor. In other words, the actual duration of the corresponding sound is $89/96 \approx 0.93$ beats. More crucially to the present study, different representations of the same note type are assigned different durations. For instance, the following three representations of a dotted quaver,



are assigned, respectively, 67, 71, and 69 units, instead of the expected $3 \times 96/4 = 72$ units. In the raw list of notes, these three representations appear as different entries, although they correspond to identical “words”. To solve this problem, the list of notes was modified semi-automatically, by manually inspecting the list to record the durations assigned to different representations of the same “word” and, then, automatically replacing the durations by their expected values (96 units for a crotchet, 72 units for a dotted quaver, and so on).

After this step, the list of notes was ready for final processing. This was done with commercial software for mathematical algebraic manipulation, Mathematica[®] 4 (www.wolfram.com), which also includes subroutines for statistical analysis and function fitting. Occurrences of different “words” were counted, rank lists were constructed, and fitting parameters were determined. A by-product of the fitting subroutine is the chi-square test of quality of fitness quoted in the main text.

Acknowledgements

I am grateful to M. A. Montemurro and I. Samengo for their critical reading of the manuscript.

References

Bernstein, L. (1973). *The Unanswered Question*. Cambridge, MA: Harvard University Press.

Boroda, M. G. and Polikarpov, A. A. (1988). The Zipf-Mandelbrot law and units of different text levels. *Musikometrika, 1*, 127--158.

Brown, C. M. and Hagoort, P. (2000). *The Neurocognition of Language*. Oxford: Oxford University Press.

Koelsch, S., Kasper, E., Sammler, D., Schulze, K., Gunter, Th., and Friederici, A.D. (2004). Music, language and meaning: Brain signatures of semantic processing. *Nature Neurosciences, 7*, 302--307.

Krumhansl, C.L. (1990). *Cognitive Foundations of Musical Pitch*. Oxford: Oxford University Press.

Lehrman, P. and Tully, T. (1993). *MIDI for the Professional*. London: Music Sales.

Lerdahl, F. and Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press.

Maess, B., Koelsch, S., Gunter T., and Friederici, A. D. (2001). Musical syntax is processed in Broca's area: an MEG study. *Nature Neurosciences, 4*, 540--545.

Manaris, B., Purewal, T., and McCormick, Ch. (2002). Progress towards recognizing and classifying beautiful music with computer. MIDI-encoded music and the Zipf-Mandelbrot law. In *Proceedings of the IEEE SoutheastCon 2002* (pp. 52--57). IEEE, CD-Rom edition. See also www.cs.cofc.edu/~manaris/publications/.

Manaris, B., Vaughan, D., Wagner, C., Romero, J., and Davis, R.B. (2003). Evolutionary music and the Zipf-Mandelbrot law: Progress towards developing fitness functions for pleasant music. In G. Raidl *et al.* (Eds.), *Applications of Evolutionary Computing* (pp. 522--534). Berlin: Springer Verlag.

Mandelbrot, B. B. (1959). A note on a class of skew distribution function. Analysis and critique of a paper by H. A. Simon, *Information and Control*, *2*, 90--99.

Montemurro, M. A. and Zanette, D. H. (2002). New perspectives on Zipf's law in linguistics: From single texts to large corpora. *Glottometrics*, *4*, 86--98.

Patel, A. D. (2003). Language, music, syntax and the brain. *Nature Neurosciences*, *6*, 674--681.

Perle, G. (1991). *Serial Composition and Atonality: An Introduction to the Music of Schoenberg, Berg, and Webern*. Berkeley: University of California Press.

Schoenberg, A. (1967). *Fundamentals of Musical Composition*. London: Faber and Faber.

Simon, H.A. (1955). On a class of skew distribution functions. *Biometrika*, 42, 425--440. Reprinted in Simon, H.A. (1957). *Models of Man. Social and Rational*. New York: John Wiley & Sons.

van Eemeren, F. H. (2001). *Crucial Concepts in Argumentation Theory*. Chicago: University of Chicago Press.

Zanette, D. H. and Montemurro, M. A. (in press). Dynamics of text generation with realistic Zipf's distribution. *Journal of Quantitative Linguistics*.

Zipf, G. K. (1935). *The Psycho-Biology of Language*. Boston: Houghton Mifflin.

Zipf, G. K. (1949). *Human Behaviour and the Principle of Least Effort*. Cambridge, MA: Addison-Wesley.

Footnotes

1. Here, I use the word *semantics* in the strict sense, namely, as the connection between symbols and the entities that they represent in the extra-symbolic world.
2. The MIDI files of the musical compositions studied in this article are available at [www.geocities.com/ benedetto_marcello/midi/](http://www.geocities.com/benedetto_marcello/midi/).

Figure captions

Figure 1: Zipf's plot (number of occurrences \underline{n} versus rank \underline{r}) for Dickens's David Copperfield. The number of different words is $\underline{V}=13,884$, and the total number of words is $\underline{T}=362,892$. In this double-logarithmic plot the straight line manifests the power-law dependence of $\underline{n}(\underline{r})$ for large \underline{r} . The dotted curve is a least-square fitting with the prediction of Simon's model, equation (1).

Figure 2: Zipf's plots for single notes in four musical compositions for keyboard. Their titles, as well as the corresponding value of \underline{V} and \underline{T} , are indicated in each panel. Curves stand for least-square fittings with the prediction of Simon's model, equation (1). The resulting exponent ν , which provides a quantitative measure of context definiteness, is given with each plot.

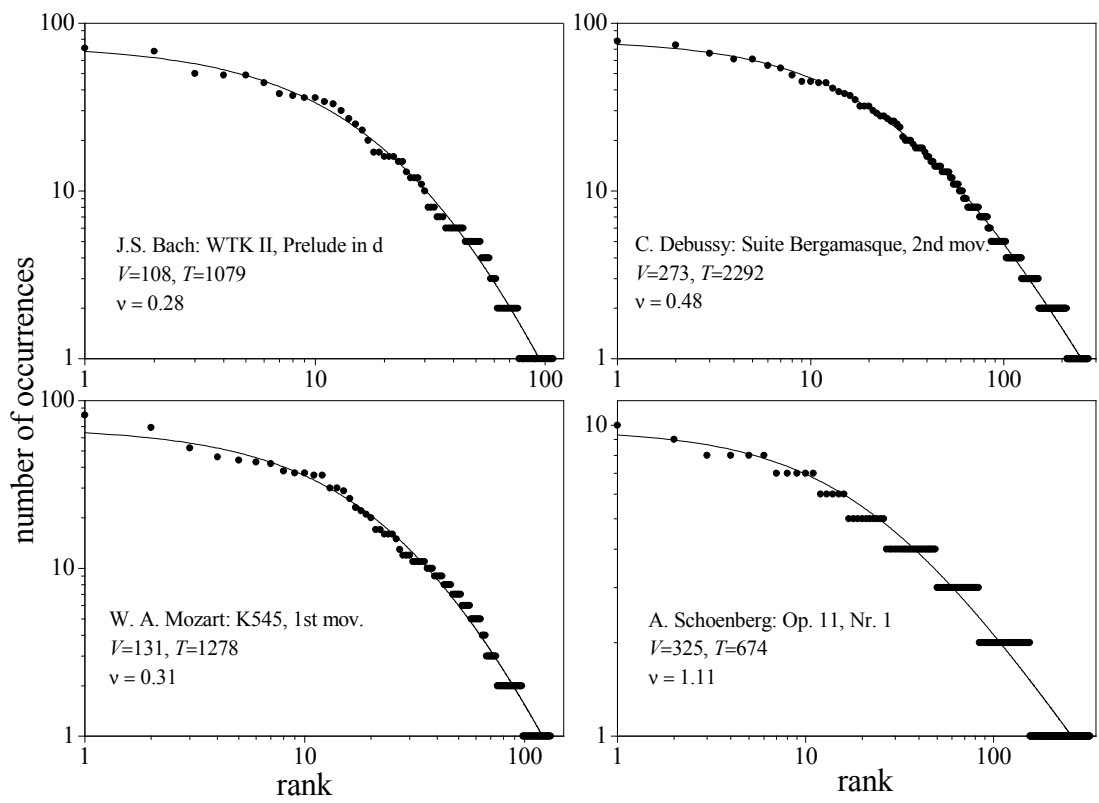


Figure 2