Computational Aesthetics and Music: the Ugly, the Small and the Boring

Stefano Kalonaris, [1] Toby Gifford, [2] Andrew Brown, [3]
[1]RIKEN AIP, Tokyo
[2]Monash University, Melbourne [3]Griffith University, Brisbane

Abstract

Contemporary musical aesthetics, as a field in the humanities, does not typically argue for the existence of aesthetic universals. However, in the field of computational creativity, universals are actively sought, with a view to codification and implementation. This article critiques some statistical and information-based methods that have been used in computational creativity, in particular their application in assessing aesthetic value of musical works, rather than the more modest claim of stylistic characterization. Standard applications of Zipf's Law and Information Rate are argued to be inadequate as computational measures of aesthetic value in musical styles where noise, repetition or stasis are valued features. We describe three of these musical expressions, each with its own aesthetic criteria, and examine several exemplary works for each. Lacking, to date, is a computational framework able to account for socio-political and historical implications of creative processes. Beyond quantitative evaluations of artistic phenomena, we argue for deeper intersections between computer science, philosophy, history and psychology of art.

1 Introduction

This paper argues that current computational methods for evaluating aesthetic value of musical pieces are too blunt to discern instances of creativity which transcend traditional notions of universal beauty. This endeavour relates to computational systems in the context of both classification/appreciation (of either human or machine artifacts) and generation. The gap between aesthetic science and aesthetic experience has been discussed in (Makin 2017) while the gap between computational aesthetics and aesthetics in the broader discourse pertaining to the philosophy of art was addressed in a previous paper (Kalonaris and Jordanous 2018). Postmodern musical works, for example, often do not conform to beauty universals and cannot be evaluated by the current methods used in computational creativity to judge aesthetic value. That is, they can be so radical and novel that they have to be considered in their own right.

It is becoming increasingly important that the field of computational creativity acknowledges the depth and complexity of musical aesthetics, and the limitations of current computational models in either generating or evaluating them. As algorithmically curated content becomes the norm, and traditional gate-keeper roles of producers/executives of record labels are outsourced to AI, it behooves us to recognize the socio-political interwinement of music.

Humans are very proficient in contextualizing creative works and in dealing with semiotic cues, while machines are notoriously restricted in this regard. Can this gap be filled? Can it help to redefine our objectives and aspirations in relation to systems able to be creative or to evaluate creative products?

Expanding the argument presented in our previous work, we provide practical examples of computational aesthetic evaluation of several pieces belonging to three musical genres that exemplify the issues faced in this endeavor. These genres are *Japanoise* (Novak 2013), *Microsound* (Thomson 2004), and *Minimal* music. They were chosen specifically because of their divergent sonic characteristics, and because they illustrate how artistic processes move beyond the expressive medium, to invoke socio-political and philosophical concerns.

By questioning traditional definitions of music and by redefining expectations for music experience, these musical genres implicitly undermine aesthetic universals. While some (Galanter 2012) persist in associating aesthetic evaluation with "normative judgments related to questions of beauty and taste in the arts" (2012, p.256), we recognize that "artworks, especially modern ones, are appreciated for other reasons besides their aesthetic qualities or beauty" (Leder and Nadal 2014, p. 445). We thus distance ourselves from anachronistic interpretations of aesthetics, and suggest that the notion of beautiful music is obsolete and irrelevant to a current discourse in aesthetics and the arts, and we hope that this can be acknowledged in the field of computational creativity.

2 Creativity and Aesthetics

Creativity and aesthetics are intrinsically bound in a feedback loop whereby the experience of a creative process, object or phenomenon is coupled with responses at an affective, sensory, and semantic level. These might affect future creative acts or the perception of new creative processes/products, and so forth. According to Rhodes (1961), creativity is a synthesis of four necessary relational modes: *product*, *person*, *process* and *press* (meaning the environment). If this true for generating art then it is likely true for evaluating art. Similar considerations can apply to aesthetics where meaning depends on context and the complex interaction of multiple factors. In its historical trajectory from the theory of taste and beauty (Hume 1757; Kant 1987) to recent discussions about post-digital art and design (Berry and Dieter 2015), discussions of aesthetics have ranged from the philosophy of fine art (Hegel 1975), philosophy of criticism (Beardsley 1981), art as experience (Dewey 1989) and pragmatist aesthetics (Shusterman 2000). It is now widely understood that an artwork does not have an innate, universal and abstract aesthetic value, but it is instead appreciated in its socio-political context. Much of this literature is yet to be acknowledged in the field of computer science (Spratt and Elgammal 2014). It has been suggested (Shusterman 1997; Bergeron and Lopes 2012) that there are three aesthetic dimensions in an experience: evaluative, phenomenological/affective and semantic. The latter, which involves a social framework that defines expectations, anticipation and preferences, is of particular importance in postmodern and contemporary expressions of art which foreground the conceptual over the perceptual. Dealing with this aspect can be problematic in the context of computational creativity, a domain which intersects both computer and cognitive sciences, philosophy, and the creative arts.

2.1 Computational Creativity

Technologies have afforded us enormous expressive and aesthetic scope, enabling novel art forms and practices to emerge. However, it seems that when considering the creativity of computational systems or their ability to judge creative products, more restrictive criteria are often employed. In large part, a system is a by-product of its designer, implicitly embedding her aesthetic/artistic objectives and preferences. To a large extent, the system's architecture (Gifford et al. 2018) will determine its creative potential. For example, a corpus-based system operating primarily in combinational and exploratory methods, will exhibit mostly p-creativity (Boden 1991). Similarly, this system's generative output will be evaluated (either by humans or by automated analysis) as a function of similarity and pertinence to the corpus it is supposed to emulate. It is difficult to envisage instances of true transformational and historical creativity (Boden 1991) for generative systems, to date.

To do that, a creative system should be able to position the art work in relation to everything that has happened before in the domain of interest (e.g., music), but also in relation to societal and political implications. This entails metacreativity (Linkola et al. 2017), or the ability to reflect, adapt and change domain knowledge via interaction with the wider environment. It might also require self-awareness, as in the ability to become the object of its own attention (Morin 2006), and a sense of identity (e.g., how to position one's self in the world). Understanding context is not limited to direct environment or the artifacts that one produced over time, but also to the environment as the cumulative effect/affect of temporal developments in multiple dimensions (possibly over a very long span, e.g., centuries), which are linked to the creative process in sometimes non-trivial ways. In general, the ability of a system to formulate an aesthetic value judgment will be hereinafter referred to as computational aesthetics (Hoenig 2005).

2.2 Computational Aesthetics

There can be two ways to consider computational aesthetics (Galanter 2012), one catering for human notions of beauty, the other catering for an emerging machine meta-aesthetics. We will focus on the approaches that have been employed to endow computational systems with the ability to evaluate the aesthetic value of both their own or of humans' creative products. Among these approaches, some have been applied to music. These aesthetic measures have been surveyed in detail in (Kalonaris and Jordanous 2018), and are based on:

- *Information/Complexity Theory* (Birkhoff 1933; Bense 1965; Gell-Mann and Lloyd 1996; Schmidhuber 2012), with music application such as the *Audio Oracle* (AO) (Dubnov, Assayag, and Cont 2007; Surges 2015)
- *Fractal Dimension and Geometry* (Spehar et al. 2003; den Heijer and Eiben 2010), which have been applied in (Voss and Clarke 1978; Manaris et al. 2007)
- *Psychology and Gestalt* (Berlyne 1960; Narmour 1990), tested and used in (Eisenberg and Thompson 2003; Brown, Gifford, and Davidson 2015)
- *Grammar*. For example, the *Generative Theory of Tonal Music* (GTTM) (Lerdahl and Jackendoff 1983), although not specifically developed for aesthetic considerations (Lerdahl and Jackendoff 1992), has been implicitly discussed in relation to aesthetics with practical computational implementations (Hamanaka, Hirata, and Tojo 2006; 2016)
- *Biology/Evolution Theory*, which include agent-based systems and have been applied to music in (Eigenfeldt, Bown, and Casey 2015; Blackwell 2007)
- *Connectionism/Neural Networks*. Systems using this paradigm have architectures particularly apt to aesthetic discrimination, albeit based on corpora similarity or reward rules. Music applications are many and increasingly ubiquitous (Eck and Schmidhuber 2002; Mehri et al. 2016; Hutchings and McCormack 2017; Hadjeres, Pachet, and Nielsen 2017)

3 Two Measures of Beauty

Aesthetic evaluation in the context of neural network-based systems includes methods often reminiscent of the Turing Test insofar as the result must be judged as sufficiently authentic/coherent with respect to a 'style' or idiom. This is a controversial approach to evaluation, which has been criticized in (Ariza 2009). While this topic is well worth exploring, we will concentrate exclusively on geometric and information theoretic methods, leaving others for future discussion. In particular, this section provides an introduction to Zipf's Law (Zipf 1949) and Information Rate, along with short musical considerations.

3.1 Zipf's Law

Zipf's Law expresses the occurrence frequency f of an n^{th} ranked event e, using the ratio

$$f(e) = \frac{1}{n^a} \tag{1}$$

with a being close to 1. This means that, for example, the second most likely event should have half the probability of the most likely event, etc.

In (Manaris, McCormick, and Purewal 2002) Zipf's Law is used as an aesthetic measure for pieces belonging to several musical styles: Baroque, Classical, several flavours of Romantic, Impressionist, 12 Tone, Jazz, Rock, Pop, Punk Rock, DNA Encoded Music, Random (white noise) and Random (pink noise). According to this study, all but the DNA and the two strands of Random exhibit a near-Zipfian distribution, with a slope of approximately -1.0 calculated over a set of descriptors based on individual pitches and intervals.

It is also claimed that Zipf's Law applied to frequency distribution of the 12 pitches in a piece "is very reliable in identifying 12-tone music, since such metric is characterized by the uniform distribution of pitches" (2002, p.5). Unfortunately this is misleading, since 12-tone music is more about the systematic organization of the 12 pitches (whereby no pitch can be re-used until the 12-tone row has been exhausted) than it is about their distribution over time. While this issue is understood and discussed by the authors of that paper, they refer to Zipf's Law outliers as *pathological* cases, and claim that "a near-Zipfian distribution is a necessary, but not sufficient condition for beautiful music" (2002, p.6).

Descriptions based on probability density calculations are clearly problematic when time is an important parameter of the phenomenon under examination. In a later paper Manaris et al. (2005) address a similar criticism of such global distribution analysis by implementing 'fractal' versions of the distribution metrics, which reflect the scaling behaviour of these distributions over different timescales. They do not, however, report on the resulting fractal dimensions encountered in their corpus; they appear to use these fractal metrics as a filter to discard pieces that are 'anomalous'.

3.2 Information Rate

According to Birkhoff (1933), the notion of beauty can be formalized by an aesthetic measure M which is the ratio between order (O) and complexity (C).

$$M = \frac{O}{C} \tag{2}$$

where $C = H(x_n)$ and $O = H(x_n) - H(x_n|x_{past})$. In these definitions, and for a given signal x, H is Shannon's entropy (Shannon 1948), x_n is the current signal, and x_{past} is the past signal. Perhaps M is a measure of orderliness rather than beauty, which does not seem to address the discourse on aesthetics that has developed in the philosophy of the art. One can note that the numerator O in Equation 2 is nearly identical to the definition of Information Rate (IR), which is used in the *Audio Oracle* (AO) (Dubnov, Assayag, and Cont 2007; Dubnov and Assayag 2012; Surges 2015) algorithm, where the entropy is substituted with a compression algorithm C which measures the number of bits required to represent the data with or without knowledge of the past. IR is defined as

$$IR(x_n, x_{past}) = C(x_n) - C(x_n|x_{past}) \tag{3}$$

The AO is a graph structure on indexed segments of a recording's audio features, which could be an alternative to spectral clustering and recurrence analysis. However, it is not used for determining structural segments, but rather to recombine sections and sub-clips of the audio analyzed, with real-time human-machine improvisation and music design as goals. The authors suggest that "the most pleasing aesthetic experiences rely on a balance between order and complexity" (Surges and Dubnov 2013, p. 3), which relate to one's ability to comprehend and appreciate a piece as a function of its novelty and repetition.

As is, IR is a measure of information and predictability that might not be appropriate in scenarios where redundancy and repetition are deliberately used for conceptual and artistic reasons. That is, repetition and redundancy have qualitatively different meanings, interpretations, implications and value in the arts, than they have in information theory or information dynamics (Tabacchi and Termini 2015). More importantly, repetition, redundancy and uniformity, are so common in postmodern music that can hardly be said to be 'pathological' (see Section 3.1). A few examples are minimal music, post-minimalism, process music, drone music, sound art, aleatoric music, free improvisation, noise, glitch, microsound, lower-case and so forth. In this respect, it is worth mentioning that in (Potter, Wiggins, and Pearce 2007) the authors applied information dynamics to the study of two works of Philip Glass. However, the focus of their study was to argue in favour of a model of analysis and segmentation of music based on human perception, rather than discussing the aesthetic implication of this method. Furthermore, and as is too often the case, this study dealt exclusively with notated musical surface. In many musical expressions that we are about to discuss the notion of symbolic representation is alien, and no analysis based on notation can be attempted.

We will focus on three of these musical/artistic expressions which seem not only to problematize the use of computational aesthetic measures, but call into question the validity of music beauty universals altogether.

4 The eye of the beholder

Music is complex and expressions of it diverse. This makes it difficult to make hard distinctions about or to isolate musical factors that identify the aesthetic value that music holds in the wider context. Furthermore, music as the object of aesthetic interest "may have very little in the way of 'facts' and almost nothing in the way of explanatory scope" (Garnett 2001, p. 23). By virtue of their pervasive nature threading through culture, politics and philosophy, the aforementioned postmodern music expressions are linked to one another, sometimes in non-obvious ways. For example, artists, composers and sound designers from one ‘genre’ might be active across several others, or some of these styles might have aesthetic and political objectives in common, although they might realize them via different means, tools and media. Detailed treatment of all of these genres exceeds the scope of this paper, but it is important to reiterate that these genres are not marginal to contemporary musical culture, nor to the discourse on aesthetics and philosophy of music. We propose a provisional categorization of these genres according to commonalities in their musical features which, it is hoped, may point toward more appropriate measures of aesthetic value in each case.

4.1 Ugly Beauty

With historical precedents perhaps traceable to Lou Reed’s *Metal Machine Music*, the Japanese take on noise started to develop in the early 1980s. Referred to as *Japanoise* in North America or as *Noise* (ノイズ) in Japan, this musical expression is “often unrelentingly harsh” (Novak 2013, p.7). Noise tunes “into the negative beauty of sublime experiences with sound” (Novak 2013, p.47). Within it, more discrimination can be made. For example the terms *liveness*, or *deadness* point at different modalities of fruition of this musical expression, whether embodied through live performance or as a listening and affective experience through recordings. These two dimensions operate as a feedback loop connecting the ephemeral of the live and the repeatability that recorded media affords. Japanoise splits musicians and listeners between those who find its membership to music inconceivable, and those who, via the distribution and circulation of recordings, help its integration into a genre of its own. Thus, “Japanoise reveals how discourses of musical globalization are continually reformed at the edge, through acts of sound-making, performance, and transcultural interpretations of popular media” (Chronus Art Center).

4.2 Small Beauty

Microsound, sometimes called *micromontage*, is the practice of using sounds “beneath the level of the note” (Roads 2004, p.vii) working between musical pitches and using small audio segments. Typically these would range between 0.1 seconds and 10 milliseconds. As a practice, it can be applied to music composition, synthesis or improvisation and live performance. Roads (2004) describes microsound techniques in detail. This musical expression has many intersections and cross-overs with *lowercase*, perhaps with more focus on the compositional aspect, rather than the improvised. Some (e.g. Phillips 2013) use both terms interchangeably, and position microsound in the broader context of *Minimalism*. Others describe it as a “cross-over space between art music, popular music, and sound art” (Haworth and Born 2018, p. 617). By virtue of its sparsity, microsound foregrounds silence and space, and is said to afford a process of self-observation, of becoming-listener. That is, microsound “elicits a unique subjectivity [. . .] and may thus be as much about the perceiving self as it is about sound and sound design” (Phillips 2006, p.233). Microsound exemplifies fundamental themes of postmodern music aesthetics such as “a complexity, nonlinearity, and multiplicity of the interpretative listening experience” (Davismoon 2016, p.269).

4.3 Boring Beauty

In this section no particular genre is referenced, although all the works chosen are in some way linked to minimalism in music.

> “And why is it called reductive, or minimal . . . ? In relation to what? Information?” (Feldman and Pellizzi 2011, p. 365)

Repetition (e.g., tape loops) and sustained tones (e.g., drones) have been a staple in the practice of notable *minimalist* composers such as La Monte Young, Steve Reich, Philip Glass and Terry Riley as much as exponents of *musiche concréte*, electroacoustic music and sonic art. For example, we can include works by Alvin Lucier, Eliane Radigue, Charlemagne Palestine and Phill Niblock, among others. It has been argued that minimalism is both an aesthetic, a style and a technique, and that “pieces focusing primarily on the process alone or pieces that lack goals and motion toward those goals best exemplify the delineation of minimalism as an aesthetic” (Johnson 1994, p.744). According to this viewpoint, we are inclined to include musical expressions such as *minimal techno*, a movement started in Detroit in the early 1990s by Robert Hood and Daniel Bell. This musical expression sports a stripped-down, essential palette of creative and musical resources (e.g., relentless repetition, understatement, sparsity, etc.) that contribute to a minimalist aesthetic.

5 Checksum

We now go back to Zipf’s Law and Information Rate and test these methods against exemplary works of the musical practices chosen in Section 4 by virtue of being indisputably influential at an aesthetic level. These can be challenging examples to deal with since they might not abide by any known tonal framework or they might foreground repetition and redundancy as their aesthetic axioms. Thus, the audio or musical features that are normally used to compute the above metrics will likely have to be edited and changed accordingly (e.g., it would not make sense to consider statistics on melodic/harmonic intervals in a buffer of noise or silence). The complete list of the musical examples considered is the following:

- Ugly Beauties
  - Hijokaidan: Self-Mutilation (from *King of Noise*, 1985)
  - C.C.C.C: Test Tube Fantasy Part One (1993)
  - Incapacitants: Libra Was Dead. Since Then, He Has Gone To Morgan Stanley (from *No Progress*, 1994)
  - Merzbow: Woodpecker No.2 (from *Pulse Demon*, 1996)
- Small Beauties
  - Bernhard Günter: Untitled I/92 (from *Un peu De Neige Salie*, 1993)
  - Ryoji Ikeda: +. (from +/-, 1996)
  - ⊘ + Noto: Untitled (from *Mikro Makro*, 1997)
  - Alva Noto: 02-10-06 Astoria 1 (from *Xerrox Vol.1*, 2007)
- Boring Beauties

- Steve Reich: Pendulum Music (1968)[1]
- Eliane Radigue: Onward 19 (from *Vice Versa*, 1974)
- Aphex Twin: Spots (from *Selected Ambient Works Volume II*, 1994)
- Robert Hood: One Touch (from *Minimal Nation*, 1994)

A few seconds of microsounds, 8 minutes of ear-aching noise or a static and motion-less drone are among the aesthetic champions chosen for this study.

5.1 Geometric Measures

To test Zipf's Law on the above works, the original features listed in (Manaris, McCormick, and Purewal 2002) (5 based on individual notes and 7 based on intervals) could not be used without incurring conceptual hurdles. Furthermore, their study used MIDI (a time-stamped representation protocol[2] for symbolic music), whereas the works considered in this study are all raw audio.

In the context of noise music, Collins (2013) has analyzed a corpus of noise music (including several Merzbow pieces) via quantitative measures in the manner of Music Information Retrieval (MIR), although with an eye to computational aesthetics. He describes six measures: perceptual loudness, transientness, sensory dissonance, spectral cresting, spectral centroid, and spectral entropy; all of which are applied to a rolling window of the audio, much as with a short-time Fourier transform (STFT), to produce a six-dimensional time series for each piece. He finds that, from a style identification perspective, these features are not sufficiently discriminatory, as they failed to separate Merzbow from the Beach Boys.

Therefore, two bespoke features were computed instead. We considered the information obtained from the constant-Q chromagram as a measure of pitch modulo 12 and the time in between detected onsets as a measure of durations. To evaluate the works chosen according to Zipf's Law we compare the results shown in Table 1 to the earlier reported statements that a $1/f$ distribution, thus a slope of (approximately) -1, is a necessary (but not sufficient) condition for beautiful music, and that examples that do not conform with this rule are to be considered pathological. The worst performing pieces are those by CCCC, Hijokaidan, Incapacitants, Merzbow, Noto, and Aphex Twin, when considering the cumulative score over both chroma and durations. Perhaps some will be pleased to see Noise fall under the not-so-beautiful music, others will interrogate whether this genre should be classed as music at all, yet some others will question the low score of the last two works (Noto and Aphex Twin). One could be even doubtful as to why *Pendulum Music* or *Onward 19* can be considered beautiful, just by virtue of having achieved an acceptable slope in one of the measures. Similarly, there could be as many objections against, or arguments in favor of, any of these works as there are listeners. Lacking, to date, a satisfactory definition of music, it is an arduous, if not futile, task to then judge what beautiful music is.

[1]In the album *SYR4: Goodbye 20th Century* (1999), by Sonic Youth

[2]https://www.midi.org/

Genre	**Track by**	**Chroma**	**Durations**
Ugly Beauties	CCCC	-0.089	0.0
	Hijokaidan	-0.494	-0.241
	Incapacitants	-0.434	-0.604
	Merzbow	-0.076	0.0
Small Beauties	Gunter	-0.269	-1.115
	⊘ + Noto	-0.514	-1.149
	Noto	-0.212	-0.283
	Ikeda	-0421	-1.062
Boring Beauties	Radigue	-1.354	0.0
	Aphex Twin	-0.306	-0.354
	Hood	-0.391	-1.317
	Reich	-1.004	-0.387

Table 1: Zipf's Law: slopes' values with respect to chroma and durations.

5.2 Information Measures

As for the Information Rate measure, we used the *PyOracle* package (Surges and Dubnov 2013), the first implementation of the AO in the Python[3] programming language.

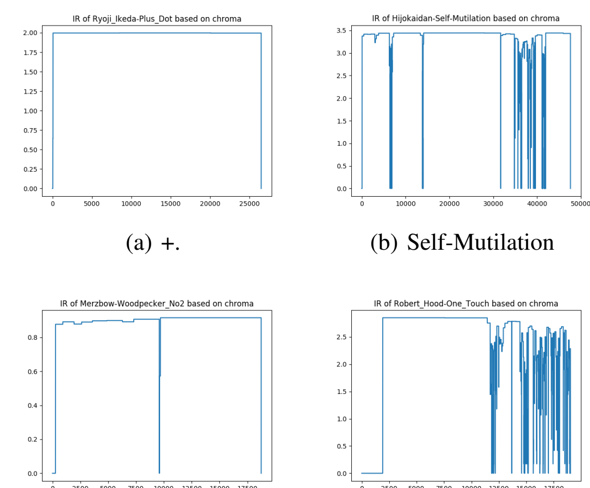

(a) +. (b) Self-Mutilation

(c) Woodpecker No.2 (d) One Touch

Figure 1: Chroma-based IR for four pieces.

In Figure 1 the IR based on chroma content of four of the works considered are shown. It can be inferred that information rate might not be a very informative measure for some of those works that have very little or no change at all, for example, Radigue's *Onward 19.* By virtue of such predictability and uniformity, and judging the aesthetic value of these pieces according to the principle of novelty vs. expectation discussed in Section 3.2, one would have to conclude that these works are not examples of beautiful music, nevertheless they are aesthetically valued by their fans.

[3]https://python.org

5.3 Timbral Measures

As we have argued, genres such as Japanoise and Microsound do not operate in the framework of 12 pitch chroma, and thus it is perhaps not surprising that analyses along these lines do not perform well as aesthetic measures. Are there timbral measures that could be more appropriate? The question of appropriate entropy/information measures for general signals is vexed. Entropy calculations are subservient to the predictive model being employed; and any form of temporal predictability renders a conventional Shannon entropy calculation (where the probabilities are taken from the global distribution) problematic. Indeed in the context of general signal analysis "the concept of complexity is far too diffuse to expect any quantitative measure of it to apply universally" (Gonzalez-Andino et al. 2000, p.49).

The approaches described above build predictive models of dimensional reductions of the signal into what are presumed to be perceptually relevant features, but do not map well to the subtle timbral evolutions that are expressed in our subject genres. As such we also analyzed our example pieces using an acoustic signal analyses borrowed from ecoacoustics; the total signal entropy (Sueur et al. 2008). This is a combination of temporal entropy and spectral entropy applied to short segments of raw audio, argued to be a proxy for the short-term complexity of the soundscape. Whilst the measure as defined can be calculated for any length of audio, it seems best suited to moment-by-moment application, resulting in a representation of the musical work as a timeline of evolving acoustic complexity.

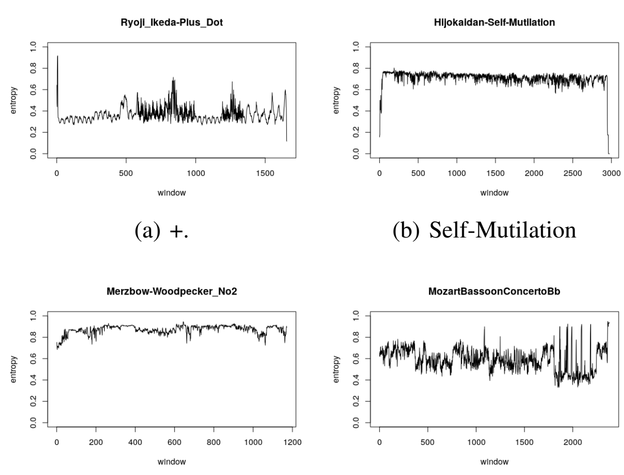

(a) +. (b) Self-Mutilation

(c) Woodpecker No.2 (d) Bassoon Concerto

Figure 2: Evolution of moment-by-moment signal entropy.

Figure 2 shows three examples of our test pieces subject to this analyses, alongside the Mozart Bassoon Concerto in Bb for comparison. The Merzbow noise piece is essentially a constant barrage of near maximal entropy. The other pieces show quite different entropy dynamics, and all seem different again to the Mozart work. We do not mean to imply that this signal entropy measure is sharp enough to capture the subtle timbral evolution and form that seems to provide interest in many noise works. Rather we hope to highlight the difficulties in formulating universal aesthetic criteria for music in general.

6 Discussion and Conclusion

Tempting as it is to seek schemata that might define a theory of experience, so that it can be reduced to a computational abstraction, the truth of the matter is that aesthetics is a complex domain which transcends formulas, generalizations and universals. On the contrary, often the approach to aesthetic experience in computer science bears similarities to that described by Goguen (2004) in the context of musical *qualia*:

> "a common approach is to 'bracket' or exile the qualitative aspects, and concentrate attention on aspects that are reducible to scientific analysis" (2004, p.5).

Epperson (1994) suggests that "musicians have much to gain from the study of aesthetics because of its profound relevance to music making" (1994, p.81). Would then, computer scientists in the business of designing creative music systems not benefit from a similar strategy? Furthermore, while we have concentrated on the gap between philosophy of art and computational aesthetics, there are other fields that would merit attentive consideration, such as the findings and developments of aesthetic theory from a psychological viewpoint. A discussion of these exceeds the scope of this paper, and it is best left for the future. In this respect, we consider the model of aesthetic appreciation proposed in (Leder et al. 2004) and updated in (Leder and Nadal 2014) as one of the most complete frameworks to date, and we endeavour to explore this further to include practical computational implementations.

We reviewed common approaches to aesthetic evaluation in the context of computational creativity with particular focus on music. We considered musical expressions that were far from being marginal, pathological or irrelevant to the artistic discourse on aesthetics and music. The results suggest these approaches have conceptual issues that are not resolvable within a purely mathematical or information-based framework. By flagging the inadequacy of some of these methods (when used in isolation) and the existing gap between aesthetic notions in philosophy of the art and psychology with those in computer science, we hope to further the discussion in the field of computational creativity, embracing the complexity of the subject matter, and of the arduous task lying ahead: the integration of a theory of experience (aesthetics) and a method of abstraction (computation).

References

Ariza, C. 2009. The Interrogator as Critic: The Turing Test and the Evaluation of Generative Music Systems. *Computer Music Journal* 33:48–70.

Beardsley, M. 1981. *Aesthetics: Problems in the Philosophy of Criticism.* Indianapolis: Hackett Publishing Company.

Bense, M. 1965. *Aesthetica Einführung in Die Neue Aesthetik.* Agis-Verlag.

Bergeron, V., and Lopes, D. M. 2012. Aesthetic theory and aesthetic science: Prospects for integration. In Palmer, S., and Shimamura, A., eds., *, with Vincent Bergeron, Aesthetic Science: Connecting Minds, Brains, and Experience*. Oxford University Press.

Berlyne, D. 1960. *Conflict, arousal, and curiosity*. McGraw-Hill series in psychology. New York, NY: McGraw-Hill.

Berry, D. M., and Dieter, M., eds. 2015. *Postdigital aesthetics: Art, computation and design*. Springer.

Birkhoff, G. D. 1933. *Aesthetic measure*. Cambridge, MA, USA: Harvard University Press.

Blackwell, T. 2007. Swarming and music. In Miranda, E. R., and Biles, J. A., eds., *Evolutionary Computer Music*. London: Springer. 194–217.

Boden, M. A. 1991. *The Creative Mind: Myths and Mechanisms*. New York, NY: Basic Books, Inc.

Brown, A. R.; Gifford, T.; and Davidson, R. 2015. Techniques for generative melodies inspired by music cognition. *Computer Music Journal* 39(1):11–26.

Chronus Art Center. Japanoise and the cultural feedback of global media. `https://bit.ly/2FheUfO`. Accessed: 2019-02-07.

Collins, N. 2013. Noise music information retrieval. In *Noise In and As Music. Huddersfield: University of Huddersfield*.

Davismoon, S. 2016. Atomisation of sound. *Contemporary Music Review* 35(2):263–274.

den Heijer, E., and Eiben, A. E. 2010. Comparing Aesthetic Measures for Evolutionary Art. In Di Chio, C.; Brabazon, A.; Di Caro, G. A.; Ebner, M.; Farooq, M.; Fink, A.; Grahl, J.; Greenfield, G.; Machado, P.; O'Neill, M.; Tarantino, E.; and Urquhart, N., eds., *Applications of Evolutionary Computation*, 311–320. Berlin, Heidelberg: Springer.

Dewey, J. 1989. *The Later Works, 1925–1953*, volume 10. Carbondale: Southern Illinois University Press.

Dubnov, S., and Assayag, G. 2012. Music Design with Audio Oracle using Information Rate. In *First International Workshop on Musical Metacreation (MUME)*, 1–1.

Dubnov, S.; Assayag, G.; and Cont, A. 2007. Audio Oracle: A New Algorithm for Fast Learning of Audio Structures. In *Proceedings of International Computer Music Conference (ICMC)*. Copenhagen, Denmark: ICMA.

Eck, D., and Schmidhuber, J. 2002. Finding temporal structure in music: Blues improvisation with LSTM recurrent networks. In *Proceedings of the 12th IEEE Workshop on Neural Networks for Signal Processing*, 747–756. Martigny, Switzerland: IEEE.

Eigenfeldt, A.; Bown, O.; and Casey, B. 2015. Collaborative Composition with Creative Systems: Reflections on the First Musebot Ensemble. In Toivonen, H.; Colton, S.; Cook, M.; and Ventura, D., eds., *Proceedings of the International Conference on Computational Creativity*, 134–141. Park City, UT: Brigham Young University.

Eisenberg, J., and Thompson, W. F. 2003. A matter of taste: Evaluating improvised music. *Creativity Research Journal* 15(2-3):287–296.

Epperson, G. 1994. Musical aesthetics: Reflections on uniting theory and practice. *Journal of Aesthetic Education* 28(1):77–82.

Feldman, M., and Pellizzi, F. 2011. A conversation on music and art: February 16, 1986. *RES: Anthropology and Aesthetics* (59/60):363–372.

Galanter, P. 2012. Computational aesthetic evaluation: Past and future. In McCormack, J., and d'Inverno, M., eds., *Computers and Creativity*. Springer Berlin Heidelberg. 255–293.

Garnett, G. E. 2001. The aesthetics of interactive computer music. *Computer Music Journal* 25(1):21–33.

Gell-Mann, M., and Lloyd, S. 1996. Information measures, effective complexity, and total information. *Complexity* 2(1):44–52.

Gifford, T.; Knotts, S.; McCormack, J.; Kalonaris, S.; Yee-King, M.; and d'Inverno, M. 2018. Computational systems for music improvisation. *Digital Creativity* 29(1):19–36.

Goguen, J. 2004. Musical qualia, context, time, and emotion. *Journal of Consciousness Studies* 11:2004.

Gonzalez-Andino, S.; Grave de Peralta Menendez, R.; Thut, G.; Spinelli, L.; Blanke, O.; Michel, C.; and Landis, T. 2000. Measuring the complexity of time series: an application to neurophysiological signals. *Human brain mapping* 11(1):46–57.

Hadjeres, G.; Pachet, F.; and Nielsen, F. 2017. DeepBach: a steerable model for Bach chorales generation. In Precup, D., and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, 1362–1371. PMLR.

Hamanaka, M.; Hirata, K.; and Tojo, S. 2006. Implementing A Generative Theory of Tonal Music. *Journal of New Music Research* 35(4):249–277.

Hamanaka, M.; Hirata, K.; and Tojo, S. 2016. *Implementing Methods for Analysing Music Based on Lerdahl and Jackendoff's Generative Theory of Tonal Music*. Cham: Springer International Publishing. 221–249.

Haworth, C., and Born, G. 2018. From Microsound to Vaporwave: Internet-Mediated Musics, Online Methods, and Genre. *Music and Letters* 98(4):601–647.

Hegel, G. W. F. 1975. *Aesthetics. Lectures on Fine Art*. Oxford: Clarendon Press. trans. T. M. Knox, 2 vols.

Hoenig, F. 2005. Defining computational aesthetics. In Neumann, L.; Sbert, M.; Gooch, B.; and Purgathofer, W., eds., *Computational Aesthetics 2005*. Girona, Spain: Eurographics Association. 13–18.

Hume, D. 1757. *Four Dissertations*. New York: St. Augustine's Press.

Hutchings, P., and McCormack, J. 2017. Using autonomous agents to improvise music compositions in real-time. In Correia, J.; Ciesielski, V.; and Liapis, A., eds., *Computational Intelligence in Music, Sound, Art and Design*, 114–127. Cham: Springer International Publishing.

Johnson, T. A. 1994. Minimalism: Aesthetic, style, or technique? *The Musical Quarterly* 78(4):742–773.

Kalonaris, S., and Jordanous, A. 2018. Computational music aesthetics: a survey and some thoughts. Dublin, Ireland:

Proceedings of the 3^{rd} conference on Computer Simulation of Musical Creativity.

Kant, E. 1987. Critique of judgment. Hackett. trans. Werner Pluhar.

Leder, H., and Nadal, M. 2014. Ten years of a model of aesthetic appreciation and aesthetic judgments : The aesthetic episode - Developments and challenges in empirical aesthetics. *British Journal of Psychology* 105(4):443–464.

Leder, H.; Belke, B.; Oeberst, A.; and Augustin, D. 2004. A model of aesthetic appreciation and aesthetic judgments. *British Journal of Psychology* 95(4):489–508.

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

Lerdahl, F., and Jackendoff, R. 1992. An overview of hierarchical structure in music. In Schwanauer, S. M., and Levitt, D. A., eds., *Machine Models of Music*. Cambridge, MA, USA: MIT Press. 289–312.

Linkola, S.; Kantosalo, A.; Männistö, T.; and Toivonen, H. 2017. Aspects of self-awareness: An anatomy of metacreative systems. In *ICCC*.

Makin, A. D. J. 2017. The gap between aesthetic science and aesthetic experience. *Journal of Consciousness Studies* 24(1-2):184–213.

Manaris, B.; Romero, J.; Machado, P.; Krehbiel, D.; Hirzel, T.; Pharr, W.; and Davis, R. B. 2005. Zipf's law, music classification, and aesthetics. *Computer Music Journal* 29(1):55–69.

Manaris, B.; Roos, P.; Machado, P.; Krehbiel, D.; Pellicoro, L.; and Romero, J. 2007. A Corpus-based Hybrid Approach to Music Analysis and Composition. In *Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 1*, AAAI'07, 839–845. AAAI Press.

Manaris, B.; McCormick, C.; and Purewal, T. 2002. Can Beautiful Music be Recognized by Computers? Nature, Music, and the Zipf-Mandelbrot Law. Technical report, CoC/CS.

Mehri, S.; Kumar, K.; Gulrajani, I.; Kumar, R.; Jain, S.; Sotelo, J.; Courville, A. C.; and Bengio, Y. 2016. SampleRNN: An unconditional end-to-end neural audio generation model. *Computer Research Repository* abs/1612.07837.

Morin, A. 2006. Levels of consciousness and self-awareness: A comparison and integration of various neurocognitive views. *Consciousness and Cognition* 15:358–371.

Narmour, E. 1990. *The Analysis and Cognition of Basic Melodic Structures: The Implication-realization Model*. Chicago, IL: University of Chicago Press.

Novak, D. 2013. *Japanoise: Music at the Edge of Circulation*. Duke University Press.

Phillips, T. 2006. Composed silence: Microsound and the quiet shock of listening. *Perspectives of New Music* 44(2):232–248.

Phillips, T. 2013. "mark the music": Microsound and becoming-silence. In *The Subject of Minimalism: On Aesthetics, Agency, and Becoming*. New York: Palgrave Macmillan US. 113–129.

Potter, K.; Wiggins, G. A.; and Pearce, M. T. 2007. Towards greater objectivity in music theory: Information-dynamic analysis of minimalist music. *Musicae Scientiae* 11(2):295–324.

Rhodes, M. 1961. An analysis of creativity. *The Phi Delta Kappan* 42(7):305–310.

Roads, C. 2004. *Microsound*. The MIT Press.

Schmidhuber, J. 2012. *A Formal Theory of Creativity to Model the Creation of Art*. Berlin, Heidelberg: Springer. 323–337.

Shannon, C. E. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27(3):379–423.

Shusterman, R. 1997. The end of aesthetic experience. *Journal of Aesthetics and Art Criticism* 55(1):29–41.

Shusterman, R. 2000. *Pragmatist Aesthetics: Living Beauty, Rethinking Art*. Rowman & Littlefield Publishers.

Spehar, B.; Clifford, C. W.; Newell, B. R.; and Taylor, R. P. 2003. Universal aesthetic of fractals. *Computers & Graphics* 27(5):813 – 820.

Spratt, E. L., and Elgammal, A. M. 2014. Computational beauty: Aesthetic judgment at the intersection of art and science. *Computing Research Repository* abs/1410.2488. URL: `https://arxiv.org/pdf/1410.2488.pdf`. Last accessed: 4th May, 2018.

Sueur, J.; Pavoine, S.; Hamerlynck, O.; and Duvail, S. 2008. Rapid acoustic survey for biodiversity appraisal. *PloS one* 3(12):e4065.

Surges, G., and Dubnov, S. 2013. Feature selection and composition using PyOracle. In *International Workshop on Musical Metacreation (MUME)*.

Surges, G. 2015. *Generative Audio Systems: Musical Applications of Time-Varying Feedback Networks and Computational Aesthetics*. Ph.D. Dissertation, University of California, San Diego.

Tabacchi, M. E., and Termini, S. 2015. Birkhoff's aesthetics, Arnheim's entropy. Some remarks on complexity and fuzzy entropy in arts. *International Journal of Computational Intelligence Systems* 8(6):1103–1115.

Thomson, P. 2004. Atoms and errors: Towards a history and aesthetics of microsound. *Organised Sound* 9(2):207–218.

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

Zipf, G. K. 1949. *Human Behaviour and the Principle of Least Effort*. Addison-Wesley.