

The theory and analysis of computer-generated music: A case-study of *Colossus*

Steven Jan *

University of Huddersfield
s.b.jan@hud.ac.uk

Abstract. While music theory and analysis have developed hand-in-hand with the historical evolution of human-generated music (“HGM”), there has been little consideration to date of how they might relate to computer-generated music (“CGM”). While CGM has been the subject of much debate as to how it might be optimised and evaluated, these have focused largely on aesthetic matters. Although music theory and analysis (for all their pretences to scientific rigour) encompass an implicit aesthetic dimension, their primary focus is on understanding *systems* of musical organisation at the extra- and intra-opus levels, respectively. Given this, they are well placed to offer insights into CGM at the level of pattern-to-pattern continuity and structural-hierarchic coherence. This paper considers a number of issues relevant to the application of music theory and analysis to CGM and, by means of a case-study of the composition *Colossus* (2010), by the *Iamus* computer, assesses how theory and analysis might contribute to the generation and evaluation of CGM, and how CGM might, conversely, motivate theory and analysis to expand its conceptual vocabulary to encompass non-human musics.

Keywords: Human-generated music (HGM), computer-generated music (CGM), theory, analysis, consciousness, memetics, computational creativity, *Iamus*, *Colossus*.

1 Introduction: Music Theory and Analysis in Historical Context

While much attention has been given to methods for *generating* music using computers (Fernández & Vico, 2013; Herremans, Chuan, & Chew, 2017); and while almost as much thought has been given to strategies for *evaluating* the outputs *as music* (Jordanous, 2012; Loughran & O’Neill, 2017), little consideration has been given, to my knowledge, to the *music-theoretical* aspects of computer-generated music (hereafter “CGM”) or to strategies for *analysing* it (but see Various, 2012 for the broader context and methodology). Nor, indeed, has the more fundamental question of developing a philosophy to determine whether such music (collectively or in individual instances) warrants – by analogy with

* I am grateful to Valerio Velardo for discussion of some of the ideas in this paper, and to the three anonymous reviewers for their helpful comments and suggestions.

human-generated music (hereafter “HGM”)¹ – such theoretical/analytical treatment in the first place.²

Music theory is one of the oldest of human endeavours, being recorded in Ancient Greece from the middle of the first millennium BCE and from a similar time in Ancient China (Christensen, 2002). In the Greek context, it was closely linked to cosmological speculation, musical intervals being regarded as sounding out mathematical truths and inscribing principles writ large in the position and movements of the planets (Clark & Rehding, 2001).³ Music analysis is a rather more recent discipline, arising tentatively in the renaissance period as an attempt to understand the nature of individual instances of music, but only fully taking shape in the early-nineteenth century. This dichotomy between theory and analysis is also one between the general and the specific, between the synchronic and the diachronic, and between the abstract and the particular, respectively.

Some common threads have linked the pursuit of these disciplines over the last two millennia, certainly in Europe. One is that they are endeavours conducted in order to understand the environment of our species, our place in the universe, and our nature: we make music in our image, so music theory and analysis have probed this simulacrum in order to glean insights into what makes us who we are. For most of our history we have believed that we are ourselves made in the image of a deity. Perhaps as a consequence, music theory and analysis have been seen as affording an insight into the mysteries of our creation and the nature of our creator. In our present atheistic and Darwinian age – if only from the perspective of educated westerners – the focus has shifted back from the divine to the human, and this introspective impulse is particularly true of recent developments, which have brought to bear insights from cognitive science (Gjerdingen, 1999, 2010) and linguistics (Patel, 2008) on the understanding of music. In this way, music is understood not as a window into our soul, or as a conduit to our creator, but as a key to unlock the complexities of the human brain and mind.

This human-centricity of music theory and analysis poses a problem for those wishing to extend it from HGM to CGM. The motivation for music-related computational creativity is (non-exclusively) binary: from a scientific perspective, the inherent complexity of music, resulting from its multiparametric combinatoriality, makes it an irresistible challenge for computer science; from a humanistic

¹ HGM comes in a dazzling variety of forms according – to give just two constraints – to the cultural background and level of training of the composer/producer. My focus here is primarily on HGM composed by trained/professional composers and written in broadly western art-music traditions.

² One should nevertheless remember that CGM is, at least partly, HGM, in the sense that the underlying generative algorithms which give rise to CGM are the product of human intelligence, albeit arguably not the specifically musical domain of that intelligence, and albeit an intelligence which – in a manner analogous, for instance, to aleatoric HGM – delegates the bulk of the decision-making to the computer.

³ Recent research suggests that electromagnetic radiation generated by stars and planets does indeed generate a “music of the spheres” (Levin, 2011), as the Greeks, and their medieval followers, believed.

perspective there is strangeness and beauty in experiencing music made by a non-human entity (when that music is not presented to listeners in the form of a Turing Test (Ariza, 2009)). This aesthetic frisson is not dissimilar to that arising from hearing the vocalizations of certain non-human animals, an issue I return to in Section 4.

Whichever motivation drives the application of music theory and analysis to CGM – the former might use it to verify the efficacy of an algorithm, the latter might use it to illuminate similar phenomena appearing in HGM – there are inevitably philosophical issues which it is timely to address. This paper attempts to review some of these issues, asking “how should music theory and analysis approach CGM, and what does CGM have to offer (‘as a goal or as a goad’ (Kerman, 1994, p. 61)) to theory and analysis?”.

To answer this question, I argue here that such music is tractable using current (and historical) approaches *to the extent to which it reflects (or appears to reflect) the operation of human-like perceptual-cognitive constraints* on its generation (Lerdahl, 1992). That is, if a piece of CGM respects the *hierarchical-grouping structure* of most HGM, then it is likely to be amenable to the same analytical methodologies which are applicable to HGM. Conversely, if a piece of CGM violates human perceptual-cognitive constraints, then – depending upon how comprehensively these constraints are violated – it is less likely to be a meaningful object for HGM-focused analytical methodologies. Yet, in this second case, it offers a challenge to theory and analysis to arrive at methodologies which are able to engage with it. Note that this distinction is not as clear-cut as the discussion in Section 2.2 implies: a piece of CGM might *mostly* adhere to human perceptual-cognitive constraints, violating them only occasionally.

By “hierarchical-grouping structure” I mean the perceptually-cognitively driven tendency of most HGM to fall into discrete groups (m(us)emes (Jan, 2007)) satisfying “Millerian” (Miller, 1956) short-term memory constraints (Deutsch, 1999); the tendency for these units to follow on from each other in coherent, quasi-teleological ways; and the tendency for this “chunking” (Snyder, 2000, pp. 53–56) to be replicated at multiple hierarchic levels, such that there exist higher-order units which, themselves, relate logically to each other in the diachronic unfolding of the music (Narmour, 1999). Most music theory and analysis has, unsurprisingly, attempted to understand music in these psychologically driven terms (Bent & Pople, 2018, sec. II), from sixteenth-century linguistically-rhetorically motivated analyses of vocal music by Burmeister; to eighteenth-century models of phrase and cadence concatenation in Koch and Kirnberger (Sisman, 1982); to Schenkerian voice-leading models; to Schoenbergian and Rétian theories of motivic transformation (Schoenberg, 1995; Réti, 1951); and, in more recent times (and perhaps going full-circle), to applications to music of Chomskyan generative-transformational linguistics (Chomsky, 1965; Lerdahl & Jackendoff, 1983; Temperley, 2001).

To explore these issues, I will undertake a limited analysis of *Colossus* (2010), a composition generated by the *Iamus* computer (Diaz-Jerez & Vico, 2017), in order to determine the extent to which the music is comprehensible using extant

theoretical frameworks and analytical methodologies (Section 3). This system relies on genetic/evolutionary algorithms as its underlying generative mechanism (Fernández & Vico, 2013, p. 546). This choice of case-study is deliberate (other systems, using other approaches, could have been selected here): because, it is hypothesised, human culture evolves in a (Universal) Darwinian fashion (cultural replicators – psychologically constrained memes – building upon genetic replicators (Dawkins, 1983; Jan, 2007)), it is reasonable to consider here CGM produced by systems drawing upon the same Variation-Replication-Selection (VRS)/evolutionary algorithm (Dennett, 1995, p. 343; Calvin, 1998, p. 21), for such music is perhaps more likely to be fundamentally akin to HGM, and thus comprehensible to a human listener (and theoretician/analyst), than that generated using other approaches.⁴ Because, however, there is no necessary connection between the operation of the VRS algorithm and the types of patterning humans find comprehensible, there is a qualification to this point, made in Section 2.2.

Before turning to consider this case-study, it is necessary to explore certain issues, five in number, which impinge upon our considerations, perhaps more philosophically than practically (Section 2).

2 Five Issues

This section considers – in less detail, owing to constraints of space, than their scope requires – the justifications for applying HGM-derived music theory and analysis to CGM (Section 2.1); the perceptual and cognitive constraints acting upon HGM (Section 2.2); the dichotomy between music as object and music as process (Section 2.3); music and/as embodiment (Section 2.4); and the relationship between the VRS algorithm and the semiological tripartition (Section 2.5).

2.1 Justification for Application

Perhaps it is first necessary to ask *whether* music theory and analysis should tackle CGM at all. There are some who would argue that, because computers lack consciousness and intentionality, any “music” they generate is but a cipher, a pale imitation of the “real” thing. Such views, as well as invoking the wider (human) debate as to what does and does not constitute music (Cassidy & Einbond, 2013), also draw on Searle’s “Chinese room” argument (Searle, 1980; Boden, 2004, pp. 289–294). This holds that the efficient implementation of well designed algorithmic processes can produce outputs which seem to evidence agency and understanding, but which are ultimately mechanistic – they are “all syntax and no semantics” (Boden, 2004, p. 290).

⁴ Beyond these constraints, however, the choice of a case-study is to some extent arbitrary. *Colossus* was selected owing to the esteem in which the *Iamus* system is held in the CGM community, the accessibility of its outputs (the system generates traditionally notated scores, not audio output), and the ready availability of HGM in a broadly similar style, reflected in the comparator piece – Schoenberg’s Klavierstück op. 11 no. 1 (1909) – discussed in Section 3.

One response to this criticism is to argue that intentionality and consciousness (Dennett, 1996; Blackmore & Troscianko, 2018), build upon such algorithmic processes rather than relying on alternatives: “their inherent causal powers give them a toehold in semantics” (Boden, 2004, p. 292), because semantics arises gradualistically from a myriad of rapidly performed syntactic operations multiply cross-referenced with numerous others to form a network of (Peircean) “interpretants” (Nattiez, 1990, p. 6). In the human brain, such syntactical interconnections have been hypothesised to underpin a Chomskyan generative/transformational “Logical Form” (LF): an assemblage of domain-specific representations into a language-syntax-mediated, domain-general structure. If the LF is tokened by sonic “lexemes” – the internally imagined and externally vocalised phonological representations of language – it becomes accessible to a “theory of mind” module (Fodor, 1983) and thereafter to consciousness itself. In this view, “perceptual and imagistic states get to be phenomenally conscious by virtue of their availability to the higher-order thoughts generated by the theory of mind system (i.e., thoughts *about* those perceptual and imagistic states)” (Carruthers, 2002, pp. 666; emphasis in the original; Jan, 2015, pp. 13–17).

Thus, while computers may not yet be capable of consciousness, they may potentially develop it by implementing such models as Carruthers’ (2002). In doing so, they would be able to ask the question “am I conscious now?” (Blackmore, 2009, p. 41) and, having reached this stage, would have the intentionality to create and appreciate music and other products of an intelligent and self-aware mind. Nevertheless, it might be argued – as discussed in Section 2.4 – that such appreciation (and thus the enjoyment that flows from it) is contingent upon physicality; and that, without the phenomenological intensity afforded by the moving and moved body, music remains as abstract as the Chinese-language instructions passed into Searle’s room.

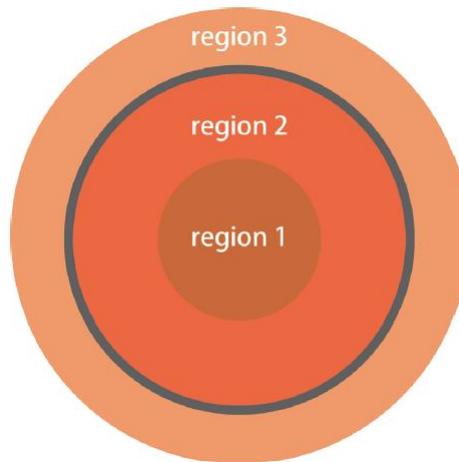
2.2 Perceptual and Cognitive Constraints

One difficulty faced by those who attempt to analyse CGM is the potential for machines (one already fully realised by some human composers) to transcend human perceptual-cognitive constraints: that is, to produce music which is beyond the complete psychological grasp of humans, such that the music is regarded, in the extreme, as noise. The “New Complexity” school, as represented by the music of Brian Ferneyhough (see, for instance, his *Mnemosyne* (1986)), arguably illustrates this tendency most clearly. In this sense, such music would occupy what Velardo terms “Region 3” of the “Circle of Sound” (Velardo, 2014).

As represented in Figure 1, after (Velardo, 2014), this circle embraces everything which might be regarded as “music” (this, of course, itself a not unproblematic concept). “Region 1” contains low-complexity music which entirely respects human perceptual-cognitive constraints. “Region 2” contains music of higher complexity than Region 1, which requires some degree of training or knowledge – implying, therefore, a “competent, experienced listener” (Meyer, 1973, p. 110) – fully to appreciate its difficulties. “Region 3” encompasses music which, on account of its violation of perceptual-cognitive constraints, is too complex for the

human mind (to cognise, if not to produce), and which might reasonably be assumed to be an inevitable product of computer, as opposed to human, creativity. Separated from Region 2 by a “horizon of intelligibility” (Velardo, 2014, p. 16), Region 3 is potentially the largest of the three, on account of its freedom from the relatively fixed and inflexible constraints operating upon human perception and cognition.

Fig. 1. The “Circle of Sound”.



Accepting this schema allows us to invoke the notions of *anthropocentric* and *non-anthropocentric* creativity (Velardo & Vallati, 2016). The former, occupying Region 1 and Region 2 of Figure 1, encompasses creativity which is by and for humans. The latter, crossing the horizon of intelligibility and occupying Region 3, encompasses creativity which is beyond human appreciation – unless we could somehow be genetically engineered in order to restructure our perceptual-cognitive apparatus – and which is therefore restricted to non-human (machine) auditors.

This distinction allows for the identification of four different types of creativity (after Velardo & Vallati, 2016, p. 6):

- Anthropocentric Creativity (Regions 1 and 2):
Humans for Humans (2H) encompasses the bulk of human creativity and its entirety before the invention of computers.
Computer-Aided for Humans (CH) relates to the use of computers as a means of augmenting human creativity.
AI for Humans (AIH) involves technology able to motivate an affirmative answer to at least the second Lovelace-question (i.e., “whether computers (now or in the future) could ever do things which at least *appear to*

be creative”) and ideally the fourth (i.e., “whether computers themselves could ever *really* be creative (as opposed to merely producing apparently creative performance whose originality is wholly due to the human programmer)”) (Boden, 2004, pp. 16-17; emphases in the original).

- Non-Anthropocentric Creativity (Region 3):

AI for AI (2AI) encompasses all creativity which is by machines and which is comprehensible only to other machines.

Non-anthropocentric creativity presupposes (i) non-anthropocentric discrimination and (ii) non-anthropocentric taste: (i) is the ability of machines to distinguish between functional uses of their competencies and artistic/aesthetic uses – akin to the ability of humans to distinguish between the skills required to solve a crossword puzzle and those required to write a sonata; and (ii) is the ability to value their creative outputs (in the light of (i)) according to various aesthetic and technical criteria.

2.3 Music as Object *versus* Music as Process

The middle of the eighteenth century – the apogee of the Age of Enlightenment – saw a profound change in the way music was understood in Europe, although one should perhaps regard this as a culmination rather than a “big-bang” phenomenon. Hitherto, music was often regarded as a *process*, and was often tied closely to its particular social function (Dahlhaus, 1983, pp. 20–23). Much music before 1700 was often performed once and then forgotten. This view is not dissimilar to the way music is (still) integrated into a majority of human cultures: it is something which forms an integral part of social functions, often being indistinguishable from dance and ritual (Merker, 2012). Gradually, music in Europe came to be seen as an aesthetic *object*, something to be preserved beyond a single performance. This was fostered by the growth in music publishing, which had taken notation (initially an *aide-mémoire*) and used it as a vehicle for dissemination and, increasingly, commercialisation of music for the new bourgeoisie of early capitalism. It was also associated with the ascendancy of the composer, increasingly specialised and separate from the performer, as not merely a craftsman, but as a genius (Stafford, 1991). The consequence of these developments was that music developed a *work concept* (Goehr, 1992), which provided a theoretical frame within which it could function as an object within a canonic discourse (Bergeron & Bohlman, 1992).⁵

The location of CGM in this dichotomy depends upon the orientation of its programmer and the target genre which is generated: the case-study considered below arises from a system designed to produce music in the image of contemporary western classical music and, as such, adopts the music-as-object model. While not a strict dichotomy, perhaps implicit in this model is the greater

⁵ While the work concept is largely restricted to European “classical” music, European popular musics, like European folk musics, share an orientation to music-as-process, as opposed to music-as-object, typical of most human musics.

priority ascribed to synchronic over diachronic factors, such that global coherence and unity are privileged. Other systems, such as Biles' *GenJam*, which is designed to generate real-time jazz improvisations (Biles, 2007), are orientated towards the music-as-process model. Perhaps implicit in this model is the greater priority ascribed to diachronic over synchronic factors, such that local (museme-to-museme) parataxis is privileged.

2.4 Music and/as Embodiment

While accepting that “Music is *fundamentally* a psychological entity, which leaves traces in the real world (audio signals and notation)” (Wiggins, Müllensiefen, & Pearce, 2010, p. 234; emphasis in the original), much music theory and analysis foregrounds music's situatedness in the physical world. Image schemata (Snyder, 2000, pp. 108, 110), arising from our nature as physical beings existing in three physical dimensions plus time, account in part for our tendency to see music in terms of notions of “up” and “down”, “fast” and “slow”, “heavy” and “light”, and related dualisms; and theory and analysis have drawn upon these metaphors (Spitzer, 2004) in such notions as the proximity of certain keys to a tonic (Schoenberg, 1983; Krumhansl & Kessler, 1982); the tension engendered by large-scale tonal structures (Rosen, 1988); the “scale-step's yearning for the tonic” (Schenker, 1980, pp. 256–257; see also Solie, 1980, p. 151); the teleology of the *Urlinie* (Schenker, 1979); the chains of implications and realizations driving melodic groupings (Narmour, 1990); and the connectedness of cognitive schemata on an imaginary thread (Gjerdingen, 2007, p. 369), to name just a few.

More profoundly, music has a visceral quality which derives in part from these image schemata: when we experience a passage as moving “up”, for instance, it “pulls” on us in an inescapably physical way. Figure 2 shows a passage from Beethoven's “Eroica” Symphony (after Liszt, 1922, p. 72), which might be experienced in these terms: we might sense a tension in our abdomen, and perhaps physically strain upwards in order to follow the perceived rising shape of the musical line and its supporting harmonies.

Of course, computers cannot (yet) experience such feelings, lacking the physical body – shaped by millennia of evolution – which gives rise to them and for which they are, ultimately, an evolutionarily driven survival mechanism. Even for quasi-autonomous robots, their “sense” of position in three-dimensional space, and their “perception” of motion, is (only) a consequence of feedback from sensors (Miranda, 2008). Of course this is essentially also the case for living organisms, but in humans, such physicality carries multiple overlays of awareness, emotion and intentionality (Section 2.1) in ways that are not (yet) available to machines.

2.5 The VRS Algorithm and the Semiological Tripartition

In terms of affecting the practice of music theory and analysis as it applies to both HGM and CGM, it is useful to consider how the VRS algorithm relates to

Fig. 2. Beethoven: Symphony no. 3 in E \flat major op. 55 (1804) "Eroica", I, bb. 182–191.

Allegro con brio

182

pp

185

cresc. *p*

188

cresc.

190

ff

the three poles of the semiological tripartition, the latter as theorised by Molino and Nattiez (Nattiez, 1990, pp. 11–12).

In Figure 3, the VRS algorithm (1) drives the *poetic* (generative) stage, by means of intra- and inter-brain memetic processes. The output product, the finished HGM (2), occupying the *neutral* (score/sound) level, is then processed via an analytical methodology (3), which is itself informed by some theoretical perspective. As such, this stage occupies the *esthesis* (reception/construction-of-meaning) pole, and gives rise to an output analysis (4), which itself may serve as the neutral level of a subsequent analytical discourse.

In Figure 4, there are two distinct poetic stages. First, the VRS algorithm drives the programming of the generative system itself (1), leading to the production of a program which is situated at the neutral level (2). This (in the case of systems using genetic/evolutionary algorithms) then invokes the algorithm to generate the output CGM (4). While this music may be analysed by a human – the principal topic of this paper – it may also be analysed by another machine (as might, of course, the music in stage 3 of Figure 3) – such as the *Humdrum Toolkit* (Huron, 2002) or the *Tonalities* software (Russ, 2004) (which are themselves products of the VRS algorithm).

As this discussion suggests, every element of these two analogous processes is made up of replicators (either in their *memotypic* (brain-stored) or their *phenotypic* (physical-world) forms (Jan, 2007, p. 30, Tab. 2.1)), sustaining the VRS algorithm at different “ontological levels” (Velardo, 2016).

3 A Case Study: *Iamus: Colossus* (2010)

A computer cluster housed in a striking tigerprint-patterned case⁶, the underlying mechanism of *Iamus* is presented under the rubric of *Melomics* (Melodic Genomics) (Sánchez-Quintana, Moreno-Arcas, Albarracín-Molina, Fernández Rodríguez, & Vico, 2013). While the technology is commercially sensitive,⁷ it is possible to understand its algorithmic basis, at least in outline, from published literature. It operates on *evo-devo* (evolutionary-developmental) principles, whereby (in biology) “evolutionary changes are interpreted as small mutations in the genome of organisms that modulate their developmental processes in complex and orchestrated ways, resulting in altered forms and novel features” (Sánchez-Quintana et al., 2013, p. 100).

One of *Iamus*’s compositions is *Colossus* (2010), for piano solo, named after the computer built to decrypt German codes during World War II by Tommy Flowers with contributions from Alan Turing. Figure 5 shows the first eight bars of the score.⁸

On first hearing, this music seems technically and stylistically convincing, having, perhaps, a flavour of the style of Messiaen (1908–1992) in its mystical and evocative textures. Cynics might argue that such a freely atonal avant-garde style is not difficult to pastiche, because musical surfaces generated by a quasi-random approach to composition (which I am not imputing to *Iamus*) may not differ markedly from those generated by strict, logical and intentional processes.⁹

The music of *Colossus* often forms patterns which are arguably coherent to a human listener, its musemes generally aligning with the perceptual-cognitive grouping criteria which govern most HGM. Moreover, there is a high degree of stylistic consistency here, with the exploration of the high registers of the piano; the use of left-hand chords which are tied across the bar line and introduced by glissandi and acciaccature; and a right-hand melody which mixes triplet and “straight” quavers. Yet the overall structure seems diffuse and lacking a clear developmental trajectory: while there is no obligation (or tradition in such a style) for an arch-shaped tension-curve, there seems no clear narrative or teleology here, as might be expected in the work of a human composer; nor is there any clear motivic development (as opposed to re-presentation), which might sustain

⁶ See <http://www.digitaljournal.com/article/327916>.

⁷ The company offers a music-streaming app, *Melomics@Life*, based on this technology (see <http://melomics.com/life>).

⁸ See also <https://goo.gl/TXin7P> for a performance with Díaz-Jerez on piano.

⁹ In a similar way, it is arguably not beyond the ability of most artistically untrained people to simulate, at least superficially, the visual style of an abstract painter like Jackson Pollock.

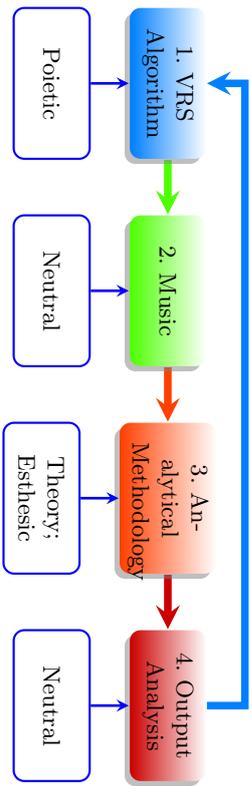


Fig. 3. Theory and Analysis in HGM.

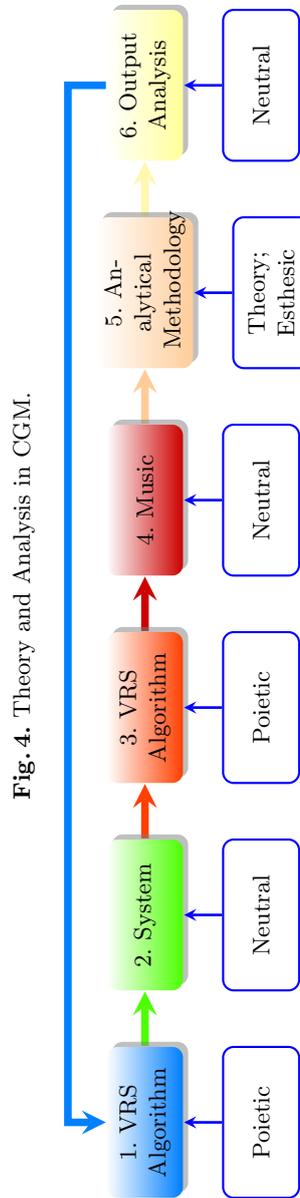


Fig. 4. Theory and Analysis in CGM.

The image displays a musical score for *Iannus: Colossus* (2010) by Iannus Computer. The score is presented in two systems, with various annotations and markings. The first system begins with a tempo marking of $\text{♩} = \text{ca. } 66$ and a dynamic of *p* (piano). It includes a *cresc.* (crescendo) marking and a triplet of eighth notes. The score is annotated with blue boxes and arrows highlighting specific musical elements and transitions. The second system features a circled measure labeled 5 (8) and various fingering numbers. The score concludes with a dynamic of *mp* (mezzo-piano) and a *cresc.* marking.

Fig. 5. *Iannus: Colossus* (2010).

Iannus Computer
(2010)

such a trajectory. In short, this piece is arguably music, but it is not particularly musical, as judged from an unavoidably biased human perspective.

An analytical methodology appropriate for attempting to understand this music is pitch-class set (PC-set) theory (Forte, 1973; Rahn, 1980), which identifies salient pitch-collections based on the chromatic set ranging from three to nine notes (of which there are 4,096) and which relates them to one of 208 fundamental “set-classes”.¹⁰ While not without its critics (McKay, 2015), PC-set theory affords the opportunity to relate seemingly disparate pitch-collections using specific set-classes (each of which has a characteristic internal interval complement or interval-class (IC) vector) as a common denominator. In HGM, two patterns of the same set-class (or, alternatively, having a “Z-relation”)¹¹ are perceived (and may have been conceived) as having a stronger synchronic/diachronic relationship than patterns without such relationships. Thus, pitch-class set correspondences may be taken as affording evidence of compositional intentionality and higher-order pitch-content planning in HGM.

As the boxes on Figure 5 suggest, and on the basis of the inevitably subjective segmentation adopted here,¹² there is, firstly, a degree of (“vertical”) recurrence of pitch-class sets evident in this extract, in the form of three appearances of 4–19 (bb. 1, 5, and 6) and two appearances of 3–11 in b. 8. Nevertheless, and secondly, while there is a degree of motivic unity engendered by the recurrent $1 \times$ quaver– $4 \times$ demisemiquaver units in bb. 1 and 3, these motives are not related by membership of a common set-class. Thirdly, there is no apparent Z-relationship between the sets identified, although alternative segmentations might reveal such relationships. Fourthly, identifying some registrally salient pitches (such as those marked by the arrows) reveals that the lower-voice pitches (D \sharp (b. 1), A \sharp (b. 4), B \natural (b. 6), and G \natural (b. 7)) spell out (“horizontally”) set-class 4–19 (no such connections are evident, however, in the upper-voice line), this set-class, as noted above, being significant vertically.

It is difficult to assess the significance of these findings, which might be purely accidental – assuming, as PC-set theory does, that in the case of HGM they are *not* accidental (i.e., that they are intentional) – in the sense of their potentially not being explicitly coded for in the algorithm.¹³ Nevertheless, the

¹⁰ The familiar major and minor triads are, in Forte’s (1973) system set-class 3–11 (the eleventh in his list of sets of cardinality three), whereas the whole-tone scale/collection is set-class 8–28. Transposition and inversion give rise to various members of each set-class, with 3–11, for instance, having 24 “distinct forms” (the twelve major and twelve minor triads).

¹¹ That is, two different set-classes which nevertheless share the same IC vector (Straus, 2005, p. 91).

¹² Segmentation is a highly controversial topic in PC-set theory (Hasty, 1981), given the arguably greater propensity of the method to confirmation bias than is the case with approaches for analysing tonal music. The segmentation of *Colossus* utilised here attempts to respect motivic and gestalt-grouping principles.

¹³ Even if such outcomes are not coded for in the algorithm – which, in the case of the VRS algorithm, cannot account for all output possibilities – this does not necessarily undermine their significance.

evo-devo algorithm underpinning *Iamus* might be responsible for the selection and replication of certain note-groups¹⁴ which, owing to their analogous interval-class content (the primary motivation, some might argue, behind Viennese early atonality, as opposed to implicit (pre-Fortean) set-class recurrence), lead to the vertical and horizontal set-class recurrences identified. That is, there may be evolutionary transmission of certain intervallic structures (in the form of PC sets functioning as unordered, interval-defined musemes) in *Iamus*'s (computer-algorithmic) implementation of cultural evolution, in a way which is analogous to the evolutionary transmission of certain biochemical structures (genes) in (nature-algorithmic) biological evolution.

While comparisons are, in some cases, odious, it is perhaps instructive to relate *Colossus* to an example of HG atonal music. Figure 6 shows the opening section of Schoenberg's (1874–1951) Klavierstück op. 11 no. 1 (1909), which we might compare with *Colossus* in terms of the four criteria just outlined: (1) recurrence of certain set-classes; (2) alignment of set-class structure with motivic structure; (3) Z-relationships between significant sets; and (4) higher-order, registrally salient (“middleground”-level) set-class structure.

The red-coloured pitches in bb. 1–2 (G♭, G♯, B♭), the blue-coloured pitches in b. 3 (D♭, E♯, F♯), the green-coloured pitches in bb. 4–5 (G♭, B♭, B♭), and the purple-coloured pitches in b. 10 (G♯, A♯, C♯) are all members of set-class 3–3 (Straus, 2005, pp. 45–47) (1), Schoenberg relating set-class membership with motivic recurrence (2) in ways not found in *Colossus*. Moreover, the bracketed melodic pitches in bb. 1–3 (upward-facing note-stems) are a member of set-class 6–Z10; and the following pitches, boxed, in the left-hand are a member of 6–Z39. These two set-classes are “Z-correspondent” with each other, sharing the IC vector 333321 (Straus, 2005, pp. 92–93) (3). Finally, the highest melodic pitches (G♭, G♯, B♭) and the lowest bass-voice pitches (G♭, G♭, B♭) (marked, respectively, by down- and up-arrows) are themselves members of set-class 3–3, forming an expression of this set-class at the middleground level (Straus, 2005, pp. 104–105) (4).

On this very limited body of evidence it would appear that, in terms of the aspects considered here (including the four specific criteria), *Colossus* possesses to some extent the hierarchical-grouping structure referred to in Section 1 and evident in the local and higher-order set-class structure of Schoenberg's op. 11 no. 1, albeit perhaps without the surface-level clarity and rigorous motivic logic of the Schoenberg piece. This potential deficiency is not necessarily to be taken as evidence that *Colossus* lacks aesthetic value, or that it is to be regarded as inferior to the Schoenberg piece. Rather, it is to acknowledge that CGM does not always end up conforming to the same perceptual-cognitive constraints as HGM despite, in this case (and from what we know of its operational principles), *Iamus* using a broadly Darwinian algorithm; and thus, in the case of *Colossus*, it may

¹⁴ It is not clear to what extent *Iamus* uses human-analogous perceptual-cognitive constraints when selecting and replicating note-patterns. Some of those in *Colossus* are, in Lerdahl's phrase, “cognitively opaque” (1992, p. 118).

potentially score less highly on rubrics deriving from analytical methodologies which have evolved to describe and explicate HGM.

4 Conclusion: Beyond Anthropocentric Theory and Analysis?

This paper has attempted to cover a good deal of ground and so, owing to the complexity of the subject, has only been able to offer some preliminary thoughts on the issues considered. As suggested in Section 1, the three main motivations for music theory and analysis are:

1. to understand the *nature of music*;
2. to understand the *nature of the human mind*; and
3. to understand a) how mind affects music; and b) whether music affects (i.e., shapes, reconfigures) mind.

In the case of their application to CGM, one might claim that the issues covered here suggest that:

1. the *nature of music* is illuminated by a comparison – if only a negative one – between HGM and CGM;
2. the *nature of the human mind* is illuminated by comparison of the human input – music-aesthetic/music-theoretical and computational – into the programming of the generative system and the configuration of the subsequent outputs;
3. a) (mind affecting music) is contingent upon 2; and b) (music affecting mind) is contingent upon the general receptivity of a human listener to CGM (assuming that listener is aware of the music’s non-human origin), and such music’s capacity (as with HGM) to expand Region 2 and to destabilise the Horizon of Intelligibility (Velardo, 2014, p. 16) (the latter perhaps acting as a long-term evolutionary selection pressure).

Given this, and in summary, one might make the following points:

- Music theory and analysis have developed alongside the music they seek to explicate, so it is not surprising that they become self-reinforcing: theory evolves to model a target which is itself constantly evolving, both owing to pressures of VRS-algorithm-driven memetic evolution; and much music, to ensure coherence, follows certain constraints of organization consolidated in theory (Nattiez, 1990, p. 135, fig. 6.1).
- While much CGM is superficially convincing in comparison to HGM, its lack of – or perceived deficiencies in – the hierarchical-grouping structure which ensures coherence (to humans) at multiple structural levels is a significant difference, often leading to a lower perceived teleological drive (Section 2.4) of CGM, in comparison with HGM (Section 3).

- This deficiency often renders CGM problematic when exposed to theoretical frameworks and analytical methodologies evolved for HGM. A tension arises owing to this: should algorithms be modified in order to generate music which is more tractable to theory and analysis (and therefore, by extension, more comprehensible to human perception and cognition); or should theory and analysis expand its own Horizons of Intelligibility, in order to accommodate the challenges of this new category of music, as it has done in the case of HGM for centuries?
- While it was assumed in Section 3 that an analyst is aware that the object of investigation is an instance of CGM, it should be acknowledged – in what might be regarded as a theoretical/analytical Turing test – that the outcome of an analysis may well be affected by knowledge of the non-human origins of CGM. If the analyst were unaware of the music’s provenance, one might (perhaps cynically) hypothesise that certain elements regarded as deficiencies in *known* CGM might be regarded as creative innovations in *assumed* HGM.

Some of these issues result from the embodied and the biological- and cultural-evolutionarily shaped nature of human music (Section 2.4), its antecedent musilanguage (Brown, 2000; Mithen, 2006), and the wider perceptual-cognitive foundations upon which these are built (Section 2.2). They currently separate HGM from CGM, but it is not inconceivable that the same evolutionary pressures might build analogous adaptations in machines; or, conversely, that evolution might reshape human perception and cognition along the lines of point (3) in the second list in this section.

Perhaps one way forward for the engagement of music theory and analysis with CGM – albeit one which implies a degree of machine subservience to, not independence from, human constraints – would be for greater use to be made of machine-learning techniques (Fernández & Vico, 2013, p. 528) in generative systems. This represents a way

to model a real human’s encounters and development with musical history. The audio corpus studies at the heart of the Music Information Retrieval (MIR) research agenda have strong potential as databases for training machine listening systems. We can envisage a future feedback loop, where output algorithmic compositions are created by systems trained on real musical examples, and algorithmic outputs may in turn become the next generation of available music. (Collins, 2018, p. 12)

A corollary to the issues considered here, and a potentially fruitful area for future research in this broad field of the analysis of non-HGM, is in the analysis of the vocalizations of non-human animals – animal-generated music, or “AGM” – which might, on some grounds, be regarded as equivalent to CGM. While animal “music” and “creativity” have been studied from a comparative perspective (Laland & Galef, 2009), and while they have been compared to human vocalizations, both musical and linguistic (Merker, 2012), they have not been studied, to my knowledge, from a music-theoretical/analytical perspective.

The assortative recombination and multi-level hierarchic structure of certain bird (Slater, 2000) and cetacean (Whitehead & Rendell, 2014) species (e.g., the

Superb lyrebird, *Menura novaehollandiae* and the Humpback whale, *Megaptera novaeangliae*) makes it at least amenable to some of the analytical methodologies discussed here. From an evolutionary perspective, some whale and bird song is at a comparable level of sophistication to CGM; and for those music-generative systems which use genetic/evolutionary algorithms, it is arguable that computers, on the one hand, and birds and cetaceans, on the other, are drawing upon similar VRS-algorithm-based processes to create their outputs, and that they have reached a comparable evolutionary stage in their development.

References

- Ariza, C. (2009). The interrogator as critic: The Turing test and the evaluation of generative music systems. *Computer Music Journal*, 33(2), 48–70.
- Bent, I. D., & Pople, A. (2018). Analysis. In *Grove music online*. Retrieved from <http://www.oxfordmusiconline.com/view/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-0000041862> doi: <https://doi.org/10.1093/gmo/9781561592630.article.41862>
- Bergeron, K., & Bohlman, P. V. (1992). *Disciplining music: Musicology and its canons* (K. Bergeron & P. V. Bohlman, Eds.). Chicago, IL: University of Chicago Press.
- Biles, J. A. (2007). Improvising with genetic algorithms: GenJam. In E. R. Miranda & J. A. Biles (Eds.), *Evolutionary computer music* (pp. 137–169). London: Springer.
- Blackmore, S. J. (2009). *Ten zen questions*. Oxford: Oneworld.
- Blackmore, S. J., & Troscianko, E. T. (2018). *Consciousness: An introduction* (3rd ed.). London: Routledge.
- Boden, M. A. (2004). *The creative mind: Myths and mechanisms* (2nd ed.). London and New York: Routledge.
- Brown, S. (2000). The “musilanguage” model of musical evolution. In N. L. Wallin, B. Merker, & S. Brown (Eds.), *The origins of music* (pp. 271–300). Cambridge, MA: MIT Press.
- Calvin, W. H. (1998). *The cerebral code: Thinking a thought in the mosaics of the mind*. Cambridge, MA: MIT Press.
- Carruthers, P. (2002). The cognitive functions of language. *Behavioral and Brain Sciences*, 25, 657–726. Retrieved from <http://dx.doi.org/10.1017/S0140525X02000122> doi: 10.1017/S0140525X02000122
- Cassidy, A., & Einbond, A. (Eds.). (2013). *Noise in and as music*. Huddersfield: University of Huddersfield Press. doi: 10.5920/noise.2013
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Christensen, T. (2002). *The Cambridge history of western music theory*. Cambridge: Cambridge University Press.
- Clark, S., & Rehding, A. (2001). *Music theory and natural order from the renaissance to the early twentieth century*. Cambridge: Cambridge University Press.

- Collins, N. (2018). The state of the artgorithm: Mass creation of algorithmic compositions. *Journal of Creative Music Systems*, 3(1).
- Dahlhaus, C. (1983). *Foundations of music history* (J. B. Robinson, Trans.). Cambridge: Cambridge University Press.
- Dawkins, R. (1983). Universal Darwinism. In D. S. Bendall (Ed.), *Evolution from molecules to men* (pp. 403–425). Cambridge: Cambridge University Press.
- Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. London: Penguin.
- Dennett, D. C. (1996). *The intentional stance*. Cambridge, MA: MIT Press.
- Deutsch, D. (1999). Grouping mechanisms in music. In D. Deutsch (Ed.), *The psychology of music* (2nd ed., pp. 299–348). San Diego, CA: Academic Press.
- Diaz-Jerez, G., & Vico, F. J. (2017). *Iamus (computer)*. Retrieved 2017-12-07, from <http://melomics.com/iamus>
- Fernández, J. D., & Vico, F. J. (2013). AI methods in algorithmic composition: A comprehensive survey. *Journal of Artificial Intelligence Research*, 48, 513–582. doi: 10.1613/jair.3908
- Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.
- Forte, A. (1973). *The structure of atonal music*. New Haven: Yale University Press.
- Gjerdingen, R. O. (1999). An experimental music theory? In N. Cook & M. Everist (Eds.), *Rethinking music* (pp. 161–170). Oxford: Oxford University Press.
- Gjerdingen, R. O. (2007). *Music in the galant style*. New York: Oxford University Press.
- Gjerdingen, R. O. (2010). Mozart's obviously corrupt minuet. *Music Analysis*, 29(1-3), 61–82. doi: 10.1111/j.1468-2249.2011.00326.x
- Goehr, L. (1992). *The imaginary museum of musical works: An essay in the philosophy of music*. Oxford: Clarendon Press.
- Hasty, C. (1981). Segmentation and process in post-tonal music. *Music Theory Spectrum*, 3, 54–73.
- Herremans, D., Chuan, C.-H., & Chew, E. (2017). A functional taxonomy of music generation systems. *ACM Computing Surveys*, 50(5), 69.
- Huron, D. (2002). Music information processing using the Humdrum Toolkit: Concepts, examples, and lessons. *Computer Music Journal*, 26(2), 11–26.
- Jan, S. B. (2007). *The memetics of music: A neo-Darwinian view of musical structure and culture*. Aldershot: Ashgate.
- Jan, S. B. (2015). From holism to compositionality: Memes and the evolution of segmentation, syntax and signification in music and language. *Language and Cognition*, 7(1), 1–38. doi: 10.1017/langcog.2015.1
- Jordanous, A. K. (2012). *Evaluating computational creativity: A standardised procedure for evaluating creative systems and its application* (DPhil, University of Sussex). Retrieved from <http://sro.sussex.ac.uk/44741/1/>

Jordanous, {_}Anna{_}Katerina.pdf

- Kerman, J. (1994). A profile for American musicology. In *Write all these down: Essays on music* (pp. 3–11). Berkeley and Los Angeles: University of California Press.
- Krumhansl, C. L., & Kessler, E. J. (1982). Tracing the dynamic changes in perceived tonal organization in a spatial representation of musical keys. *Psychological Review*, 89(4), 334–368.
- Laland, K. N., & Galef, B. G. (2009). *The question of animal culture*. Cambridge, MA: Harvard University Press.
- Lerdahl, F. (1992). Cognitive constraints on compositional systems. *Contemporary Music Review*, 6(2), 97–121. Retrieved from <http://www.tandfonline.com/doi/abs/10.1080/07494469200640161> doi: 10.1080/07494469200640161
- Lerdahl, F., & Jackendoff, R. (1983). *A generative theory of tonal music*. Cambridge, MA: MIT Press.
- Levin, J. (2011). *The sound the universe makes*. Retrieved from https://www.ted.com/talks/janna{_}levin{_}the{_}sound{_}the{_}universe{_}makes/transcript
- Liszt, F. (1922). *Franz Liszts musikalische Werke. Bearbeitungen Band III: L. van Beethoven, Symphonien nr. 1-5* (J. Vianna da Motta, Ed.). Leipzig: Breitkopf und Härtel.
- Loughran, R., & O'Neill, M. (2017). Limitations from assumptions in generative music evaluation. *Journal of Creative Music Systems*, 2(1). Retrieved from <http://jcms.org.uk/issues/Vol2Issue1/limitations-from-assumptions/article.html>
- McKay, J. Z. (2015). The problem of improbability in music analysis. In M. Ayari, J.-M. Bardez, & X. Hascher (Eds.), *Proceedings of the conference l'analyse musicale aujourd'hui: crise ou (r)évolution?, université de strasbourg, 19–21 november 2009* (pp. 77–90). Strasbourg: Editions Delatour.
- Merker, B. (2012). The vocal learning constellation: Imitation, ritual culture, encephalization. In N. Bannan (Ed.), *Music, language, and human evolution* (pp. 215–260). Oxford: Oxford University Press.
- Meyer, L. B. (1973). *Explaining music: Essays and explorations*. Chicago: University of Chicago Press.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97.
- Miranda, E. R. (2008). Emergent songs by social robots. *Journal of Experimental and Theoretical Artificial Intelligence*, 20(4), 319–334.
- Mithen, S. (2006). *The singing neanderthals: The origins of music, language, mind and body*. London: Weidenfeld & Nicolson.
- Narmour, E. (1990). *The analysis and cognition of basic melodic structures: The implication-realization model*. Chicago: University of Chicago Press.
- Narmour, E. (1999). Hierarchical expectation and musical style. In D. Deutsch

- (Ed.), *The psychology of music* (2nd ed., pp. 441–472). San Diego, CA: Academic Press.
- Nattiez, J.-J. (1990). *Music and discourse: Toward a semiology of music* (C. Abbate, Trans.). Princeton, NJ: Princeton University Press.
- Patel, A. D. (2008). *Music, language, and the brain*. New York: Oxford University Press.
- Rahn, J. (1980). *Basic atonal theory*. New York and London: Longman.
- Réti, R. (1951). *The thematic process in music*. London: Faber.
- Rosen, C. (1988). *Sonata forms* (2nd ed.). New York: Norton.
- Russ, M. (2004). “Fishing in the right place”: Analytical examples from the Tonality project. *Music Analysis*, 23(2-3), 195–244.
- Sánchez-Quintana, C., Moreno-Arcas, F., Albarracín-Molina, D., Fernández Rodríguez, J. D., & Vico, F. J. (2013). Melomics: A case-study of AI in Spain. *AI Magazine*, 34(3), 99–103.
- Schenker, H. (1979). *Free composition* (E. Oster, Ed.). New York: Longman.
- Schenker, H. (1980). *Harmony* (O. Jonas, Ed. & E. Mann Borgese, Trans.). Chicago: University of Chicago Press.
- Schoenberg, A. (1983). *Structural functions of harmony* (L. Stein, Ed.). London: Faber.
- Schoenberg, A. (1995). *The musical idea and the logic, technique and art of its presentation* (P. Carpenter, S. Neff, & S. A. C. Neff Patricia, Eds.). New York: Columbia University Press.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–424. doi: <https://doi.org/10.1017/S0140525X00005756>
- Sisman, E. (1982). Small and expanded forms: Koch’s model and Haydn’s music. *Musical Quarterly*, 68, 444–475.
- Slater, P. J. B. (2000). Birdsong repertoires: Their origins and use. In N. L. Wallin, B. Merker, & S. Brown (Eds.), *The origins of music* (pp. 49–63). Cambridge, MA: MIT Press.
- Snyder, B. (2000). *Music and memory: An introduction*. Cambridge, MA: MIT Press.
- Solie, R. A. (1980). The living work: Organicism and musical analysis. *Nineteenth-Century Music*, 4(2), 147–156.
- Spitzer, M. (2004). *Metaphor and musical thought*. Chicago: University of Chicago Press.
- Stafford, W. (1991). *The Mozart myths: A critical reassessment*. Stanford, CA: Stanford University Press.
- Straus, J. N. (2005). *Introduction to post-tonal theory* (3rd ed.). Upper Saddle River, NJ: Prentice-Hall.
- Temperley, D. (2001). *The cognition of basic musical structures*. Cambridge, MA: MIT Press.
- Various. (2012). Special issue: Mathematical and computational approaches to music: Three methodological reflections. *Journal of Mathematics and Music*, 6(2). doi: <https://doi.org/10.1080/17459737.2012.712354>

- Velardo, V. (2014). The sound/music dilemma: Why is it that all music is sound but only some sounds are music? In *The sound ambiguity conference*.
- Velardo, V. (2016). Recursive ontology: A systemic theory of reality. *Axiomathes*, 26(1), 89–114. doi: 10.1007/s10516-015-9272-0
- Velardo, V., & Vallati, M. (2016). Beyond anthropocentric creativity: A general framework for describing creative systems. In *International conference on tools with artificial intelligence*. San Jose, CA.
- Whitehead, H., & Rendell, L. (2014). *The cultural lives of whales and dolphins*. Chicago: University of Chicago Press.
- Wiggins, G. A., Müllensiefen, D., & Pearce, M. T. (2010). On the non-existence of music: Why music theory is a figment of the imagination. *Musicae Scientiae, Discussion*, 231–255.