Universal Acid in the Computer Chip: Music, Memetics and Metacreation

Steven Jan*
University of Huddersfield
Queensgate
Huddersfield HD1 3DH
s.b.jan@hud.ac.uk

Abstract

Universal Darwinism (UD) (Plotkin, 1995) holds that the "evolutionary algorithm" (Dennett, 1995, pp. 50–52) operates across the interconnected realms of a "recursive ontology" (Velardo, 2016) that binds together all that exists. Indeed, UD maintains that all phenomena in the universe are emergent properties of Darwinian processes of variation, replication and selection. If true, an evolutionary view of culture must take priority over more "creationist" accounts that rely upon the mysterious intercession of inspiration and imagination and the supposed conscious agency of the composer (or programmer). Accepting the logic of UD it follows that computer-generated music (CGM) must in some senses be explainable by evolutionary precepts. The most widely accepted (and critiqued) theory of the operation of UD in culture, memetics, arguably has significant explanatory power for human-generated music (HGM). This paper explores how AIMC can be understood in its light, even when the outputs of generative systems seem far removed from the structural norms, aesthetic values and sound-worlds, of HGM.

Keywords: Evolution, Universal Darwinism (UD), memetics, human-generated music (HGM), computer-generated music (CGM).

1 Introduction: The Power of the VRS Algorithm

The grand project initiated by Darwin (2008) in the nineteenth century took a further leap in the 1980s and 1990s when Dawkins (1983b) and Plotkin (1995), after earlier thinkers, theorised the extension of evolutionary thought, via the notion of Universal Darwinism (UD), to other realms. They held that the mechanism of variation, replication and selection (Dennett, 1995, p. 343; Calvin, 1998, p. 21) – the "VRS algorithm" – was agnostic as to substrate, seizing hold of any patterning in any domain that had the capacity to cause copies to be made of itself, and, from those seemingly unpromising scraps of proto-information, building up systems of immense complexity. Thus, and irrespective of the "ontological category" – physical, biological, psychological or socio-cultural – in which it operated (Velardo, 2016, p. 104), the VRS algorithm served as a "universal acid" (Dennett, 1995, p. 63), dissolving away rival theories and overcoming all resistance in order to reshape everything in its path.

When applied to music, as indeed to any other substrate, it is necessary to identify the *units of selection* (Lewontin, 1970) – the patterns that are subject, as the input and output configurations, to the VRS algorithm. By analogy with genes, those patterns in music that form perceptually/conceptually discrete, bounded particles – short, demarcated sequences of pitches with a particular rhythmic garb, or *vice versa* – are most likely to serve as the "optimon", or the "the unit of natural selection" (Dawkins, 1983a, p. 81). Such patterns – memes (Blackmore, 1999) or, in my coinage, *musemes* (Jan, 2007) – exist fundamentally as brain-stored constellations of neuronal interconnection (Calvin, 1998), which one might term the *memome*, by analogy with the *genome* of biological replication. Their real-world manifestations, or *phemotype*, by analogy with *phenotype* of biology, are the behaviours and artefacts that brain-stored m(us)emes motivate, such as symbolic representations in musical

*https://pure.hud.ac.uk/en/persons/steven-jan. I am grateful to Robin Laney for his helpful comments on an earlier version of this paper, and to the anonymous reviewers for their feedback.

notation, the actions of performers, the crystallisation of those actions in various storage media (Jan, 2007, 30, Tab. 2.1), and the motivation to make computer programs that make music.

The reason that (foreground-level)[2] musemes are described above as "short, demarcated sequences" is that they are forced to be thus by the constraints of human perception and cognition, itself a survival-related aptation (Gould & Vrba, 1982) of our species driven by the VRS algorithm. This mental apparatus forms a sieve or filter, one that tightly restricts what we can retain in short-term memory. Miller's "magical number seven, plus or minus two" (1956) captures an important truth about the human capacity for (indeed our handicap of) "chunking" (Snyder, 2000, pp. 53–56); and the gestalt principles of similarity, proximity, and common direction, as formalised in Narmour's implication-realisation theory (1990, 1992), impose further limitations upon what is and is not "cognitively opaque" for our species (Lerdahl, 1992).

Machines are, of course, not intrinsically subject to such constraints: their RAM is not restricted to *c*. 7-byte chunks; and their speed of (serial) processing outstrips that of humans, even if they lack our massively complex (parallel, consciousness-generating) neural architecture (Dennett, 1993; Dennett, 2017). Yet because the field of music generation has hitherto primarily aimed to create facsimiles of human-generated music (HGM) in computer-generated music (CGM),[3] some systems implicitly adopt human constraints if they employ certain algorithms, notably machine-learning or expert-knowledge approaches. If such systems are thus grounded on HGM, then they will inevitably assimilate, implicitly or explicitly, some or all of the constraints that shaped that HGM, even though they are not themselves subject to those constraints: they are hobbled by the limitations that apply to humans but that do not apply to themselves (if one accepts this anthropomorphism). Thus, and if the "blind watchmaker" (Dawkins, 2006) of the VRS algorithm is given due credit, then HGM is music that musemes caused to be made using human brains; and CGM is music that musemes caused to be made using the computers that human brains caused (via memes) to be made.

Yet some CGM is capable, in principle, of breaking free of those constraints – as are, of course, certain genres of HGM, such as in the case of the integral serialism of Milton Babbitt or the New Complexity of Brian Ferneyhough – generating music that, to a human, is beyond their comprehension and/or aesthetic appreciation. In Velardo's terms, this would constitute an example of *non-anthropocentric creativity*, specifically "AI for AI (2AI)", whereby what is generated is not comprehensible to a human but is so to another machine (Velardo, 2014; Velardo & Vallati, 2016).[4] Of course the notion of comprehensibility – let alone the question of aesthetic appreciation – in relation to machines is highly problematic, because it goes beyond perception (input) to encompass cognition and even consciousness. The existence in principle of 2AI CGM would, nevertheless, appear to suggest that such systems, and the music they generate, are seemingly resistant to Dennett's Universal Acid – that they do not replicate musemes after all.

This paper argues that this is not the case, and that the VRS algorithm is indeed operative in all CGM as it is in all HGM. It makes this argument in two ways: (i) by expanding the concept of a museme to encompass patterns that are cognitively opaque to a human but that are not so to a machine; and (ii) by theorising the operation of the VRS algorithm not just in *generating the output CGM* (in the "mind" of the system) but also the VRS algorithm's operation in *generating the generative system itself* (in the mind(s) of the system's programmer(s)). To these ends, section 2, apropos point (i) above, explores how the nature of generative algorithms affects the perceptibility and comprehensibility of the output music to human listeners, and to notional machine listeners. Section 3, apropos point (ii), considers how the VRS algorithm is implicated in the genesis and development of generative systems. Section 4 outlines some further implications of a memetic orientation in relation to the ontology of HGM and CGM, and considers the nature of authenticity in the latter.

[2] Section 2 admits the existence of more abstract structures, at the middleground and background levels (Schenker, 1979), as musemes, these categories also encompassing replicated formal-structural types.

[3] The veracity of some of these facsimiles has been assessed by their use as experimental materials in Turing Tests (Cope, 2001, p. 21; Ariza, 2009).

[4] Anthropocentric creativity, by contrast, encompasses: (i) that by one human that is accessible to another human ("2H" creativity); (ii) computer-aided creativity for humans ("CH", whether "transhuman" or "posthuman", in Yenidogan's (2021) sense); or (iii) that of an AI that is comprehensible to a human ("AIH" = most CGM).

2 Algorithms for CGM Constrained by HGM and Those Unconstrained

As suggested in Section 1, it is in the nature of certain music-generative algorithm-classes that they will emulate the structure – both high- and low-level – of HGM, and they will therefore produce CGM that is organised by means of the recursive-hierarchical structure-generation via allele-parataxis ("RHSGAP") that is hypothesised to operate in HGM (Jan, 2022, sec. 3.5.2).[5] The RHSGAP model holds that HGM arises from the linear concatenation, or parataxis, of musemes to form a *musememeplex*, these in turn generating higher-order structures at what might be regarded as a Schenkerian middleground level (Schenker, 1979). If the same museme/plex(-allele)-sequence recurs in two or more contexts, the same higher-order structure, a *musemesatz* (a replicated structural plan, defined by common periodicity and melodic and/or tonal-harmonic shape), will be generated/replicated (Jan, 2010). In their implementation of RHSGAP, certain CGM systems produce "artificially" that which arises "naturally" in culture. Whether this means that one should still refer to such patterning as constituted of musemes – with all the implications of human-centredness that are associated with this term – is a moot point. Alternatively, one could adopt Blackmore's notion of the *treme* (2015), meaning a replicated particle native to some digital environment.[6] Nevertheless, this consideration of nomenclature should not detract from the acknowledgement that, whether one calls such entities memes, musemes or tremes, they are replicated particles selfishly "trying" to survive in a ruthless competition with their alleles (Distin, 2005).

Capable of sustaining a counterfactual track of musemic evolution, those algorithms most likely to be constrained to the RHSGAP-driven organisational principles of HGM, and thus to the world of 2H, CH and AIH creativity, include both knowledge/rule-based systems and machine-learning systems, *provided they are orientated, by programming or training, to HGM*.[7] The former class of algorithms are generally encoded, top-down, with information abstracted by a human from HGM, and so will tend to result in the generation of output CGM that aligns with the input HGM corpus from which the rules are derived. The latter class of algorithms, specifically neural networks, abstract patterning bottom-up from an input HGM corpus, resulting in the replication of that patterning in the output CGM. A third class of algorithms, optimisation systems, will tend to achieve similar outcomes to the first two classes via different methods. A genetic/evolutionary algorithm, particularly if it is trained on an input HGM corpus, will model the VRS algorithm in the operation of its code; and it will tend to alight upon patterns in its output CGM that align with the constraints of HGM musemes and their recursive-hierarchical structuring, because its fitness function is coded by a human (or, as in the case of MacCallum, Leroi et al. (2012) and MacCallum, Mauch et al. (2012), is enacted by humans) in the light of HGM. The same algorithms can potentially break these constraints, and thus explore the immense terrain of 2AI creativity, *provided they are decoupled, by programming or training, from HGM*.

The former (HGM-constrained) group are considered in Section 2.1, the latter (HGM-unconstrained) group in Section 2.2.

2.1 Algorithms Constrained by HGM

Taking an example from Sturm's *folk-rnn* system (Sturm et al., 2016; Sturm & Ben-Tal, 2017; see also Sturm & Maruri-Aguilar, 2021), Fig. 1a shows an example of one of the system's training inputs, the jig "Thank God we're surrounded by water" (melody no. 2611, second version, in the online folk-music community *The Session*'s (Various, 2021) database); Fig. 1b shows one of the system's generated outputs (melody no. 2857, as transcribed in Sturm, 2017b, p. 2871); and Fig. 1c shows

[5] By analogy with a gene allele, an allele x^1 of a museme x is one that has a broadly analogous form to x and that can, by virtue of this similarity, occupy the same functional-structural node(s) or *locus*/*loci* in a composition as x. Examples include, but are certainly not limited to, the schemata of galant music (Gjerdingen, 2007), each of whose variant forms can be deployed at specific structural *loci* to assemble phrases.

[6] Apropos the memome-phemotype distinction, the fundamental (memomic) form of a treme is a sequence of electronic impulses in a digital system. While a treme might, from a human-centric perspective, be regarded as the phemotype of a m(us)eme, understood from its own perspective it is a replicator in its own right, whose phemotypic products include the light waves emanating from computer monitors and the sound waves produced by speakers.

[7] See the taxonomy of CGM systems/algorithms in Fernández and Vico (2013) (adapted in Jan (2022, sec. 6.5)). See also Herremans et al. (2017) and Herremans (2022).

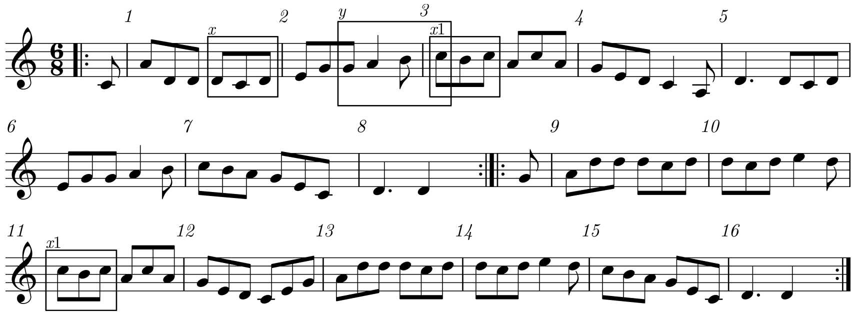

(a) Example of *folk-rnn*'s Training Input: Melody no. 2611, "Thank God we're surrounded by water", Second Version.

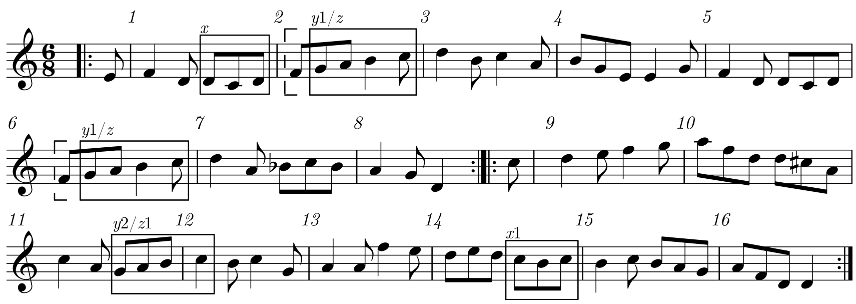

(b) Example of *folk-rnn*'s Generated Output: Melody no. 2857, Original.

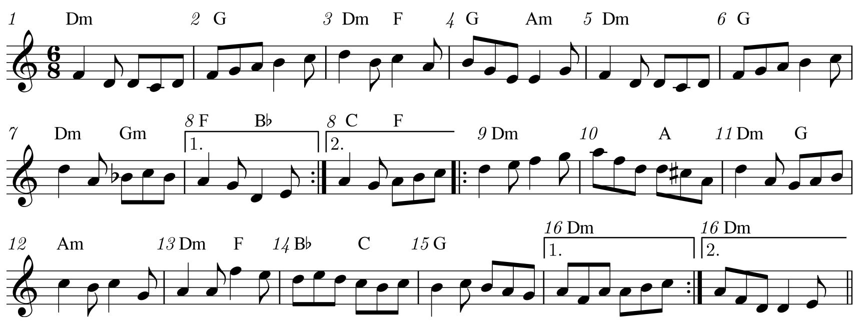

(c) Example of *folk-rnn*'s Generated Output: Melody no. 2857, Version 2, as Modified by Sturm and by Jan.

Figure 1: Examples of *folk-rnn*'s Training Input HGM and Generated Output CGM.

an improved version of the melody of Fig. 1b, with suggested harmonisation, by Sturm (2017a), transposed to the "Ddor" of the original version in Sturm (2017b, p. 2871).[8]

Given the nature of this repertoire, and the constraints of its 6/8 metre, its constituent musemes tend either to be bar-long units (usually demarcated by barlines), or half-bar units (often three-quaver groups) that form components of one or more bar-long units, and thus are replicated independently of the encompassing bar-level grouping. At a higher structural-hierarchic level, several melodies may be connected by a shared musemesatz, whereby the same foreground-level musemes (or their alleles) generate a common middleground-level structure. To take just one of several examples of foreground-level musemes, and while there is no bar-length museme in Fig. 1a that recurs in Fig. 1b,[9] there are several recurrent half-bar-length patterns linking the two melodies. For instance, the lower-auxiliary-note pattern d^1–c^1–d^1 of b. $1^{4–6}$ (boxed) in Fig. 1a – museme *x*, a candidate "antecedent coindex" or posited source in the training input – recurs in the same locus of Fig. 1b – a candidate "consequent coindex" or posited derivative in the output melody. An allele of this

[8] I have made a few further modifications to Sturm's improved version in Fig. 1c, correcting some odd harmonisations and, most significantly, changing the c^2 of b. $11^{1–2}$ to d^2.

[9] This is not to say that one or more of the bar-length musemes in Fig. 1b do not have homologues in other melodies in the training dataset.

museme (sharing its ↓–↑ contour but not its intervallic structure), c^2–b^1–c^2, museme x^1, occurs as an antecedent coindex in bb. $3^{1–3}$ and $11^{1–3}$ of Fig. 1a and as a consequent coindex in b. $14^{4–6}$ of Fig. 1b.[10]

More abstractly, alleles of the four-note pitch-sequence g^1–a^1–b^1–c^2 of bb. 2^3–3^1 of Fig. 1a – museme *y*, a common tetrachord in tonal music – appear in bb. $2^{2–6}$ and $6^{2–6}$ (y^1), as well as in bb. 11^4–12^2 (y^2), of Fig. 1b. In y^1, however, the g^1–c^2 tetrachord is a subset of a longer pattern, f^1–g^1–a^1–b^1–c^2[11], hence its designation (as the superscript indicates) as an allele, not a (unmodified) consequent coindex, of *y*. y^1 could also be regarded as a separate museme, *z*, as was suggested in note 10 apropos museme x^1, this reading being more credible given that y^1, if the prefatory (dotted-boxed) f^1 is included, is a pentachord, not (like *y*) a tetrachord. Moreover, *y* as read here arguably possesses more coherence and perceptual-cognitive salience if regarded as the longer pattern e^1–g^1–g^1–a^1–b^1–c^2 (bb. 2^1–3^1).

To generalise the issues raised by this second constellation of examples, while there may be a pitch-sequence recurrence in an output that aligns with the same pitch-sequence in the input, certain overlaps may not be sufficiently salient – owing to issues related to their alignment with the rhythmic-metrical structure and their pitch/harmonic-content – to warrant a secure judgement of memetic replication were they to be found in HGM. Of course, *folk-rnn* is not directly constrained by human perceptual-cognitive limitations, and is presumably sensitive to any pitch-sequence recurrences, irrespective of their metrical placement. The latter situation might arise "accidentally" in HGM owing to the concatenation of musemes – via what might be termed "overlapping" coindexation – and then subsequently be foregrounded by metrical and rhythmic factors.

2.2 Algorithms Unconstrained by HGM

Writing of the imperceptibility of the serial organisation of Boulez's chamber cantata *Le Marteau sans Maître* (1954), Lerdahl (1992, p. 97) notes that "[t]here is a huge gap here between compositional system and cognized result". Even allowing for the effect of the repeated hearings required to comprehend complex music, Lerdahl concludes that

> competent listeners to *Le Marteau*, even after many hearings, still cannot even begin to hear its serial organization. For many passages they cannot even tell if wrong pitches or rhythms have been played. The piece is hard to learn by ear in a specific sense; its details have a somewhat statistical quality. Conditioning, in short, does not suffice. For another thing, *Le Marteau* does not feel structurally complex in the way, for example, that compositions by Beethoven or Schoenberg do. Vast numbers of non-redundant events fly by, but the effect is of a smooth sheen of pretty sounds. The listener's processing capacities, in short, are not overwhelmed. (Lerdahl, 1992, p. 98)

The same description might equally apply to certain instances of CGM, including those produced by the *Iamus* computer (Diaz-Jerez, 2011; Sánchez-Quintana et al., 2013), whose programmers specifically aim for it to generate music in broad alignment with some of the concerns of the most progressive "classical" musics of the second half of the twentieth century. An instance of its output, a short but representative extract from the composition *Hello World!* (2011), for violin, clarinet and piano, is shown in Fig. 2.

While the underlying algorithm of *Iamus* is not described in detail in the literature, owing in part to its commercial sensitivity – see, however, Puy, 2017, sec. 2 – its basic approach is genetic-algorithmic. Specifically, its operation is described as based on "evo-devo" – evolutionary-developmental – principles of gene-mutation (memome/tremome in memetic terms) and embryological development (phemotype-generation) (Sánchez-Quintana et al., 2013, p. 100). While it is not clear how the biological relationship between the genotype and the phenotype is mapped to their posited cultural equivalents in the form of the memome and phemotype in this system, it is evident that the evolutionary processes involved generate and mutate pitch and rhythm sequences of considerable complexity; indeed, in Lerdahl's (1992, p. 118) term, they are often of high cognitive opacity. For example, the

[10] Of course, museme x^1 could be regarded as a museme in its own right, in which case our conception of its ontology, and thus its symbology, would change. More fundamental is the issue of the point at which a form of a m(us)eme is better regarded as a mutation rather than as an allele, and thus potentially as the start of a new line of evolutionary development

[11] Sturm's indicated "G" chord in Fig. 1c implies the f^1 is the seventh of this chord and that the a^1 and c^1 are passing notes.

Figure 2: *Iamus* Computer: *Hello World!* (2011), bb. 60–63.

grouping structure of Fig. 2 is difficult to parse, partly on account of the cross-instrument three-against-two rhythmic dissonance. Atonal (but not serial), "its details", to re-quote Lerdahl, "have a somewhat statistical quality", seemingly lacking the tight correlation between motivic development and interval/pitch-class-set structure (Forte, 1973) idiomatic in much free-atonal music from Schoenberg and Bartok onwards. While perhaps partly a matter of enculturation and taste, and while well beyond the scale of the pieces in Fig. 1, it is difficult to discern a sense of large-scale structure in *Hello World!*, and indeed in other outputs by *Iamus*, such as the solo piano piece *Colossus* (2010) (Jan, 2018). The lack of motivic and tonal-harmonic development gives the music an aimlessness which, while perhaps to some extent "intentional" – a feature not a bug – is arguably un-typical of the best HGM in the style the system appears to be aiming to replicate.

Thus, Fig. 2 appears to pave the way towards CGM of increasing inaccessibility to humans, the same degree of inaccessibility which, some would argue, has provoked a schism between listeners and composers of HGM in the western art-music tradition (Rosen, 2001). As a potential driver of this human-CGM schism, Collins (2018, pp. 11–12) "envisage[s] 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 [for system training]". This loop – which implies that implicated machines will gradually purge music of its HGM-ness and concomitantly intensify its CGM-ness – opens up the possibility of a machine generating CGM that replicates patterning in a manner quite unlike that in HGM. Specifically, one might conceive – as is possible in Fig. 2 – of "distributed" patterns being formed by non-contiguous pitch and rhythmic structures connected together by virtue of their co-replication. While this happens in HGM when a middleground-/background-level musemeplex/musemesatz is replicated across two or more works, its equivalent in CGM might well seem "flat" to humans – it would, in Lerdahl's phrase, "not feel structurally complex" – but it would not have the same flatness, perhaps, to the machine. Moreover, while it is the interstitial (foreground-level) patterning and not the generated (middleground- and background-level) structural nodes that are more salient in HGM for humans, the converse might potentially be true in CGM for machines.

3 Applying the VRS Algorithm to the Generation of Generative Algorithms

To formalise how the VRS algorithm might drive not just the generation of HGM and CGM but also the *genesis of generative systems themselves*, Fig. 3 represents the intersection of the three components of Molino's semiological tripartition – the poietic (production), neutral (product), and esthesic (reception) levels, in terms of which the ontology of a work of art can be framed Nattiez (1990, pp. 11–12) – with the VRS algorithm (Jan, 2018). Note that Nattiez is concerned with the dialogue between music and a "metalanguage" that describes it (Fig. 3a) – some analytical exegesis underpinned by a theoretical framework – and this serves as a starting point for an expansion of the Molino/Nattiez model to encompass the memetics (musical and verbal-conceptual) of HGM (Fig. 3b) and of CGM (Fig. 3c).

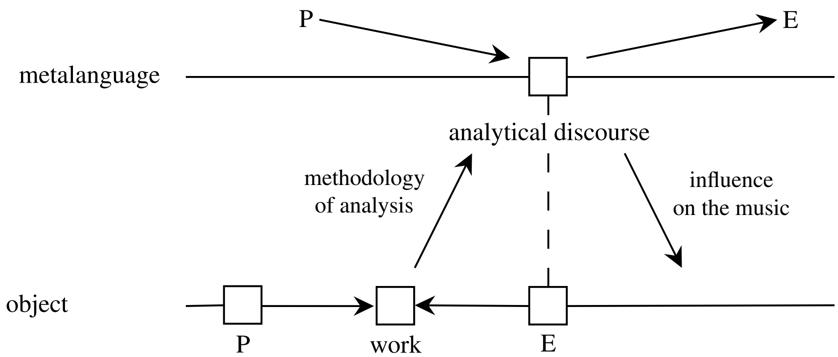

(a) Nattiez on Object, Metalanguage and Method.

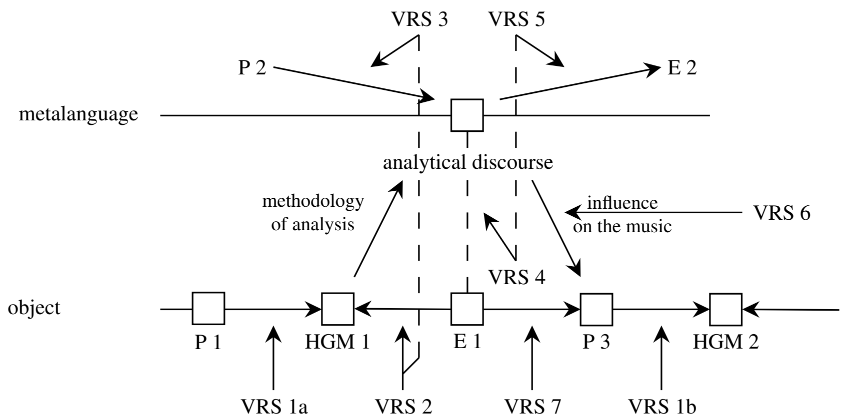

(b) HGM.

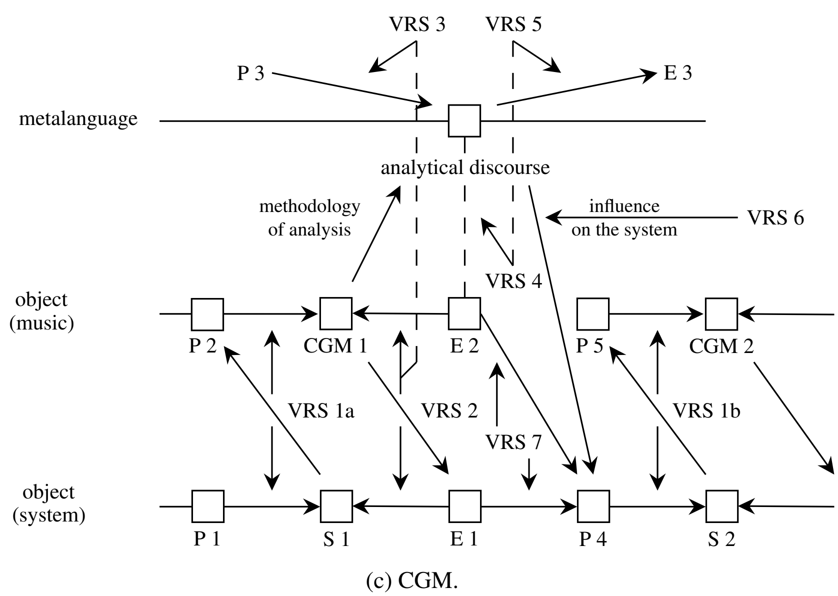

(c) CGM.

Figure 3: Generative Processes in HGM and CGM. *Key*: P = poiesis/poietic; N = neutral; E = esthesis/esthesic; S = system; VRS = Variation-Replication-Selection algorithm.

3.1 VRS, HGM and Critical-Scholarly Discourse

Fig. 3a (Nattiez, 1990, 135, Fig. 6.1) represents the relationship between the poiesis and esthesis of a work and the "analytical discourses" – indeed any scholarly perspective on music, acting as a metalanguage – that mediate these two processes. In brief, a work motivates (and validates) a certain "methodology of analysis", the latter forming part of a music-scholarly discourse associated with that work and with other similar works. This discourse feeds back into the poietic process via its "influence on the music": a way of analysing one piece, underpinned by a particular theoretical standpoint, motivates another composer to write a later piece in the light of that theoretical standpoint. Key to this model is the notion that a theoretical/analytical discourse is *itself* the product of poiesis, and is also subject to esthesis, these processes operating in tandem with those governing the work of art that the discourse attempts to model or elucidate.

Fig. 3b adapts Fig. 3a to incorporate the action of the VRS algorithm on poiesis and esthesis. That is, it sees the production of music and its reception, and the production and reception of theoretical/analytical discourses, as driven by the three interconnected evolutionary processes of variation, replication and selection. The making of a work ("HGM1", for instance), is shown as arising from the action of the VRS algorithm on musemes in the brain of the composer or improviser ("VRS1"), which are varied, differentially selected according to some criterion of fitness, and replicated by virtue of their incorporation into a phemotype represented by the neutral level – the physical form or "trace" – of HGM1. Again, as with Fig. 3a, the theoretical/analytical discourse arises from the same evolutionary processes that govern the aetiology of the work; and the musemes and verbal-conceptual memes (the latter constituting the theoretical/analytical discourse) engage in *coevolutionary* interaction in order to further their replicative advantage. That is, what benefits a museme may also benefit a verbal-conceptual meme, and in this circumstance it is therefore to both of their long-term survival advantage to evolve ever closer structural and conceptual alignment.

3.2 VRS, CGM and Critical-Scholarly Discourse

Fig. 3c is an extension of Fig. 3a and Fig. 3b in that, to the evolution of music and to the evolution of theoretical/analytical discourses on music, it adds the *evolution of generative systems* that produce, in this case, CGM. In other words, it wraps up the processes leading to the creation of a music-generative system – the poiesis of the system, its reification at the neutral level, and its esthesis by other programmers – into the same Universal-Darwinian epistemology as those processes governing the generation of music and those processes governing the generation of ideas about music. To take just one segment of Fig. 3c by way of illustration, P4 (the poietic process that leads to the second iteration, or variant (S2), of the generative system) is informed by the response (E1) to the original version of the system (S1); by the response (E2) to the music (CGM1) that system-version produced; and by the influence (VRS6) on P4 by the theoretical/analytical discourse that fed upon CGM1. In more abstract terms, these interconnected processes have leveraged coevolutionary interaction between memomic (brain- and machine-stored) tremes, musemes and memes encoding programming languages, music and words, respectively, in order to generate various phemotypic products – on-screen and machine-language code, internalised music-sound patterns and internalised word-sound patterns, respectively – that inscribe evidence of their interaction and interdependence. From the perspective of selfish survival, the computer-code tremes owe their survival, in part, to the nature of the CGM musemes to which they give rise, and to the benevolence of the verbal-conceptual memeplexes that mediate the reception of the CGM and thus, indirectly, the generating system's code-tremes.

4 Aura and Authenticity in CGM

Beyond the ability of humans to perceive and cognise the outer extremes of CGM's patterning – beyond what Velardo (2014, 16, Fig. 1) terms the "horizon of unintelligibility – perhaps what ultimately separates HGM from CGM is the former's possession, and the latter's lack, of what Walter Benjamin termed (in 1935) "aura". In his view, "that which withers in the age of mechanical reproduction is the aura of the work of art", because "the technique of reproduction detaches the reproduced object from the domain of tradition" (2015, p. 215). As with that of natural objects, the aura of historical art-objects is "the unique phenomenon of a distance, however close it may be"

(2015, p. 216).[12] While Benjamin was concerned with the mechanical *reproduction* of the work of art, CGM involves its digital *production*, often to create a facsimile (a reproduction) of a stylistic original. In Benjamin's formulation, the aura arises from the distant materiality of the art-object – the historically removed reality of the *Mona Lisa*, for instance, *versus* the contemporarily proximate artificiality of its rendition on postcards in the Louvre's gift shop.[13]

In the distinction between HGM and CGM, any aura possessed by the latter arguably inheres – or is felt by a human to inhere – in its capacity to ape the former, and to emulate the former's "domain of tradition", captured imperfectly by CGM's training datasets. Yet CGM is not produced directly by an ape; it merely possesses, second-hand, aspects of our primate physiology, our embodiment (Shapiro, 2011; Cox, 2016), and our m(us)eme-driven consciousness (Dennett, 2017, ch. 14). Indeed, there is a direct connection between our human physicality, the image schemata this physicality engenders (arising from living in three dimensions on a planet with gravity; Snyder, 2000, pp. 108, 110), and the nature of musemes: when a musical pattern rises, we feel ourselves impelled upwards by it; when a pattern falls, we imagine the effect of its weight upon our body, predicting viscerally the point of impact. These alignments between musical-pattern contour and embodiment are fitness-enhancing tricks deployed by selfish musemes; but they rely upon sensations that no machine can yet experience and, thus lacking, the machine cannot directly suffuse a work of CGM with aura, understood in this post/neo-Benjaminian sense as inhering in the intersection of our human physicality with the temporality of our cultural context and inheritance. The aura inherent in HGM is nevertheless preserved, to some extent, via the various human contributions to CGM's algorithms, configuration and input data (Fig. 3c) and the resultant moulding of their outputs in our perceptual-cognitive-embodied image.

Via the property of aura an art-object acquires a patina of authenticity. In an understandably human-centric model, Moore (2002) articulates three categories of authenticity (orientated around, but generalisable from, popular music), aligned with the three persons of grammar. In his formulation,

> *authenticity of expression*, or … 'first person authenticity', arises when an originator (composer, performer) succeeds in conveying the impression that his/her utterance is one of integrity, that it represents an attempt to communicate in an unmediated form with an audience.… 'second person' authenticity, or *authenticity of experience*, … occurs when a performance succeeds in conveying the impression to a listener that that listener's experience of life is being validated, that the music is 'telling it like it is' for them.… *authenticity of execution*, or … 'third person authenticity'[,] … arises when a performer succeeds in conveying the impression of accurately representing the ideas of another, embedded within a tradition of performance. (2002, 214, 220, 218, emphases in the original)

Third-person authenticity does not apply directly to the dichotomy between HGM and CGM, unless the system is intended to reflect (apropos the domain of art-processes outlined in note 12) the nuances of specific performance styles (Miranda et al., 2010; Costalonga et al., 2019) – or if, more radically, "composition" or "improvisation" is substituted for "performance". In machine-for-human (AIH) contexts, and assuming the generative system is taken to be the "originator", in Moore's sense, then first-person authenticity would arise if the human listener perceived the "integrity" of the machine's utterance (the CGM). Second-person authenticity would arise if that machine were able to convince a human listener that the CGM resonates in some way with their corporeal experience, perhaps by deploying musemes that encode our physicality. While these authenticities are grounded on a human-memetic (2H) basis – m(us)emes are the common currency of aura and authenticity for our species – it is possible to envisage their operation in purely machine-tremic (2AI) contexts. Beyond issues of machine perception and cognition, let alone consciousness and taste, in a putative "first-machine authenticity", one machine would compute the integrity, perhaps via blockchain-based methods, of another, as expressed through its creative outputs; in "second-machine authenticity", the creative utterances of one machine would strike another as resonating with, and "validating", its own *in silico* "experience".

[12] The reference here to art-*objects* reminds us of the complementary domain of art-*processes* (see Taruskin, 1995). In CGM, the latter include those simulated by real-time improvisatory systems, some – such as Biles' *GenJam* (Biles, 2007) – acting in (CH) collaboration with human partners.

[13] This is a distinction that is inherently problematic with music, given a work's non-identity with its score.

5 Conclusion

This paper has outlined several ways in which Universal Darwinism impinges upon the theory and practice of CGM. Firstly, it has stressed that the nature of the generative algorithm has a significant bearing upon the alignment between the output CGM and (directly or indirectly) the input HGM, and with it the (to humans) perceptual-cognitive salience and coherence of the CGM (Section 2). Secondly, it has argued that the VRS algorithm impinges upon the machine generation of music at a number of levels: not just at the level of the system producing the CGM, but also at that of the aetiology, via cultural evolution, of the generative system itself (Section 3). Thirdly, it has suggested that the auratic quality of HGM is often depauperate in CGM, because its tremes may lose the visceral connection between pattern and affect intrinsic to HGM (Section 4). These issues both connect and separate HGM and CGM, and they underpin questions related to the nature, purpose and value of CGM. Nevertheless, in this aetiological and ontological separation between the two musics – and despite their common underpinning by the VRS algorithm – lies the potential for the radical, post-human evolution of music. Do androids dream of electric sheep (Dick, 1968)? They might come to do so, but they might also dream of, and indeed consciously bring into being, electric sonatas, whose tremes they have replicated from the cultural traditions of their human creators. Moving beyond what we might understand music to be, they might also conjure up other, as-yet-unknown electric sound-worlds – audacious digital fantasias of unimaginable complexity and richness – of which their distant human creators could not possibly conceive or perceive.

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.

Benjamin, W. (2015). The work of art in the age of mechanical reproduction (H. Zorn, Trans.). In H. Arendt (Ed.), *Illuminations* (pp. 211–244). The Bodley Head.

Biles, J. A. (2007). Improvising with genetic algorithms: *Genjam*. In E. R. Miranda & J. A. Biles (Eds.), *Evolutionary computer music* (pp. 137–169). Springer.

Blackmore, S. J. (1999). *The meme machine*. Oxford University Press.

Blackmore, S. J. (2015). The next replicator. In J. Brockman (Ed.), *What to think about machines that think: Today's leading thinkers on the age of machine intelligence* (pp. 150–152). Harper.

Calvin, W. H. (1998). *The cerebral code: Thinking a thought in the mosaics of the mind*. MIT Press.

Collins, N. (2018). "… there is no reason why it should ever stop": Large-scale algorithmic composition. *Journal of Creative Music Systems*, *3*(1). https://doi.org/10.5920/jcms.525

Cope, D. (2001). *Virtual music: Computer synthesis of musical style*. MIT Press.

Costalonga, L. L., Pimenta, M. S., & Miranda, E. R. (2019). Understanding biomechanical constraints for modelling expressive performance:a guitar case study. *Journal of New Music Research*, *48*(4), 331–351.

Cox, A. (2016). *Music and embodied cognition: Listening, moving, feeling, and thinking*. Indiana University Press.

Darwin, C. (2008). *On the origin of species by means of natural selection, or the preservation of favoured races in the struggle for life* (G. Beer, Ed.). Oxford University Press.

Dawkins, R. (1983a). *The extended phenotype: The long reach of the gene*. Oxford University Press.

Dawkins, R. (1983b). Universal Darwinism. In D. S. Bendall (Ed.), *Evolution from molecules to men* (pp. 403–425). Cambridge University Press.

Dawkins, R. (2006). *The blind watchmaker: Why the evidence of evolution reveals a universe without design* (Thirtieth anniversary). Penguin.

Dennett, D. C. (1993). *Consciousness explained*. Penguin.

Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. Penguin.

Dennett, D. C. (2017). *From bacteria to Bach and back: The evolution of minds*. Allen Lane.

Diaz-Jerez, G. (2011). Composing with melomics: Delving into the computational world for musical inspiration. *Leonardo Music Journal*, *21*, 13–14. https://doi.org/10.1162/LMJ_a_00053

Dick, P. K. (1968). *Do androids dream of electric sheep?* Doubleday.

Distin, K. (2005). *The selfish meme: A critical reassessment*. Cambridge University Press.

Fernández, J. D., & Vico, F. J. (2013). AI methods in algorithmic composition: A comprehensive survey. *Journal of Artificial Intelligence Research*, *48*, 513–582. https://doi.org/10.1613/jair.3908

Forte, A. (1973). *The structure of atonal music*. Yale University Press.

Gjerdingen, R. O. (2007). *Music in the galant style*. Oxford University Press.

Gould, S. J., & Vrba, E. S. (1982). Exaptation: A missing term in the science of form. *Paleobiology*, *8*(1), 4–15.

Herremans, D. (2022). *The computer-generated music repository*. Retrieved February 12, 2022, from http://dorienherremans.com/cogemur/

Herremans, D., Chuan, C.-H., & Chew, E. (2017). A functional taxonomy of music generation systems. *ACM Computing Surveys*, *50*(5), 1–33. https://doi.org/10.1145/3108242

Jan, S. B. (2007). *The memetics of music: A neo-Darwinian view of musical structure and culture*. Ashgate. https://doi.org/10.4324/9781315085951

Jan, S. B. (2010). Memesatz contra Ursatz: Memetic perspectives on the aetiology and evolution of musical structure. *Musicae Scientiae*, *14*(1), 3–50. https://doi.org/10.1177/102986491001400101

Jan, S. B. (2018). The theory and analysis of computer-generated music: A case-study of *colossus*. In R. Loughran (Ed.), *Proceedings of the third conference on the computer simulation of musical creativity, University College Dublin, 20–22 August 2018*. University College Dublin. https://csmc2017.wordpress.com/

Jan, S. B. (2022). *Music in evolution and evolution in music*. Open Book Publishers. https://doi.org/10.11647/OBP.0301

Lerdahl, F. (1992). Cognitive constraints on compositional systems. *Contemporary Music Review*, *6*(2), 97–121. https://doi.org/10.1080/07494469200640161

Lewontin, R. C. (1970). The units of selection. *Annual Review of Ecology and Systematics*, *1*, 1–18.

MacCallum, R. M., Leroi, A. M., Mauch, M., Welburn, S., & Bussey, C. (2012). *DarwinTunes: Survival of the funkiest*. http://darwintunes.org/

MacCallum, R. M., Mauch, M., Burt, A., & Leroi, A. M. (2012). Evolution of music by public choice. *Proceedings of the National Academy of Sciences*, *109*(30), 12081–12086. https://doi.org/10.1073/pnas.1203182109

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., Kirke, A. J., & Zhang, Q. (2010). Artificial evolution of expressive performance of music: An imitative multi-agent systems approach. *Computer Music Journal*, *34*(1), 80–96. https://doi.org/10.1162/comj.2010.34.1.80

Moore, A. (2002). Authenticity as authentication. *Popular Music*, *21*(2), 209–223.

Narmour, E. (1990). *The analysis and cognition of basic melodic structures: The implication-realization model*. University of Chicago Press.

Narmour, E. (1992). *The analysis and cognition of melodic complexity: The implication-realization model*. University of Chicago Press.

Nattiez, J.-J. (1990). *Music and discourse: Toward a semiology of music* (C. Abbate, Trans.). Princeton University Press.

Plotkin, H. C. (1995). *Darwin machines and the nature of knowledge: Concerning adaptations, instinct and the evolution of intelligence*. Penguin.

Puy, N. G.-C. (2017). On the ontological category of computer-generated music scores. *Journal of Creative Music Systems*, *1*(2). http://jcms.org.uk/issues/Vol1Issue2/on-the-ontological-category/on-the-ontological-category.html

Rosen, C. (2001). *Critical entertainments: Music old and new*. Harvard University Press.

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 (Der freie Satz)* (E. Oster, Ed. & Trans.). Longman.

Shapiro, L. A. (2011). *Embodied cognition*. Routledge.

Snyder, B. (2000). *Music and memory: An introduction*. MIT Press.

Sturm, B. L. T. (2017a). *folk-rnn (v2) tune #2857*. Retrieved August 6, 2021, from https://highnoongmt.wordpress.com/2017/12/02/folk-rnn-v2-tune-2857/

Sturm, B. L. T. (2017b). *folk-rnn session books, volumes 1–10, and v3 volumes 1–4*. Retrieved August 6, 2021, from https://highnoongmt.wordpress.com/2017/08/10/folk-rnn-session-books-volumes-1-10-and-v3-volumes-1-4/

Sturm, B. L. T., & Ben-Tal, O. (2017). Taking the models back to music practice: Evaluating generative transcription models built using deep learning. *Journal of Creative Music Systems*, *2*(1). http://jcms.org.uk/issues/Vol2Issue1/taking-models-back-to-music-practice/article.html

Sturm, B. L. T., & Maruri-Aguilar, H. (2021). The Ai music generation challenge 2020: Double jigs in the style of *o'neill's 1001*. *Journal of Creative Music Systems*, *5*. https://www.jcms.org.uk/article/id/950/

Sturm, B. L. T., Santos, J. F., Ben-Tal, O., & Korshunova, I. (2016). Music transcription modelling and composition using deep learning. In V. Velardo (Ed.), *Proceedings of the first conference on computer simulation of musical creativity, University of Huddersfield, 17–19 June 2016* (pp. 1–16). University of Huddersfield. https://csmc2016.wordpress.com/proceedings/

Taruskin, R. (1995). *Text and act: Essays on music and performance*. Oxford University Press.

Various. (2021). *The session*. Retrieved August 6, 2021, from https://thesession.org/

Velardo, V. (2014). The sound/music dilemma: Why is it that all music is sound but only some sounds are music? *Proceedings of the sound ambiguity conference*. http://ambiguity.amuz.wroc.pl

Velardo, V. (2016). Recursive ontology: A systemic theory of reality. *Axiomathes*, *26*(1), 89–114. https://doi.org/10.1007/s10516-015-9272-0

Velardo, V., & Vallati, M. (2016). *A general framework for describing creative agents*. arXiv: arXiv: 1604.04096. https://arxiv.org/abs/1604.04096

Yenidogan, B. (2021). How to talk about AI art and music: An onto-ethico-epistemological debate between transhumanism and posthumanism. In A.-M. Gioti & G. Eckel (Eds.), *Proceedings of the second conference on AI music creativity, University of Music and Performing Arts, Graz, 18–22 July 2021*. Institute of Electronic Music; Acoustics (IEM), University of Music; Performing Arts. https://aimc2021.iem.at/wp-content/uploads/2021/06/AIMC_2021_Yenidogan.pdf