Defining Inspiration?

Modelling non-conscious creative process

Geraint A. Wiggins

What composers do

The title of this collection, 'The Act of Musical Composition: Studies in the Creative Process', might be taken to carry with it some suppositions about the nature of possible answers to the questions it implies. The purpose of this chapter is to deconstruct some aspects of those suppositions and to tease apart the tangled and inscrutable network of happenings that constitute the construction of a piece of music by a musician. The narrative here is intended to propose a hypothesis accounting for human musical creativity, which is a large-scale endeavour. For that reason, I do not present detail of the computational mechanisms on which I rely for evidence, nor of the empirical work done to validate them as cognitive models. The interested reader is invited to learn about the detail in the various published papers that I cite along the way.

My colleagues and I have argued elsewhere[1,2] that music needs to be studied as a primarily psychological construct, for it is from the psychology that music's equally

---

[1]Geraint A. Wiggins, 'Semantic Gap?? Schemantic Schmap!! Methodological considerations in the scientific study of music', in *Proceedings of 11th IEEE International Symposium on Multimedia*, (IEEE, 2009), pp. 477–482.

important social, sociological and shared aesthetic aspects emerge. This requirement, I claim, holds no less of composition than of any other aspect of musical behaviour; indeed, it is easiest to identify in the Romantic notion of the composer struggling alone in his[3] artistic garret, for he is abstracted from all social context other than that encoded in his memory, in his chosen notation, and perhaps that implicit in the design of his instrument, if he uses one. My approach, therefore, begins from psychology.

Composers are often placed on metaphorical pedestals, even by the most extraordinarily gifted instrumentalists, at least in the non-Popular music world. As a composer myself, I have often been asked, ‘How can you *do* that?’ But the only answer I can give, which is no answer at all, is ‘How can you *not* do that?’, by which I mean that to generate new pieces of music is so fundamentally a part of my nature that I can no more imagine not doing it than I can imagine not tasting food or not feeling the keyboard on which I am currently typing. The British composer Richard Rodney Bennett concurs: ‘I didn’t ever decide I was going to be a composer. It was like being tall. It’s what I was. It’s what I did.’[4] Of course, there is a more detailed answer, which says, ‘I take my ideas and then work them through into finished pieces of music, using such-and-such an approach’, and I could discuss the nature of the ideas, which might be motivic, timbral, structural, metaphorical, or any combination

---

[2]Geraint A. Wiggins, Daniel Müllensiefen, and Marcus T. Pearce, ‘On the non-existence of music: Why music theory is a figment of the imagination’, *Musicae Scientiae*, *Discussion Forum 5* (ESCOM, 2010), pp. 231-255.

[3]Sic. Romantic notions are rarely gender-neutral.

[4]Nick Wroe, ‘A life in music: Richard Rodney Bennett’, *The Guardian* (London, 22nd July, 2011).

of these. My point here is that, for me, at least, the generation of musical ideas, of one kind or another, is on-going, involuntary, and fundamental to my being. It follows logically that many such ideas lie abandoned, forgotten amidst the turmoil of everyday existence, though I cannot be sure of this, because it is I who have forgotten them.

One supposition implied by the current volume's title might be that there is a single identifiable Act of musical creation, which is part of a single, universal – 'the' – Creative Process. According with that position, there are myths about Wolfgang Amadeus Mozart being able to 'see' the entirety of a composition in one creative flash. However, these stories are probably derived from an inaccurate précis of Mozart's own introspective description:

> When I am, as it were, completely myself, entirely alone, and of good cheer – say traveling in a carriage, or walking after a good meal, or during the night when I cannot sleep; it is on such occasions that my ideas flow best and most abundantly. Whence and how they come, I know not; nor can I force them. Those ideas that please me I retain in memory, and am accustomed, as I have been told, to hum them to myself.
>
> All this fires my soul, and provided I am not disturbed, my subject enlarges itself, becomes methodized and defined, and the whole, though it be long, stands almost completed and finished in my mind, so that I can survey it, like a fine picture or a beautiful statue, at a glance. Nor do I hear in my imagination the parts successively, but I hear them, as it were, all at once. What a delight this is I cannot tell! All this inventing, this producing takes place in a pleasing lively dream. Still the actual hearing of the

> toutensemble is after all the best. What has been thus produced I do not easily forget, and this is perhaps the best gift I have my Divine Maker to thank for.[5]

In this account, there is no single creative flash, but rather the spontaneous emergence of initial ideas over a period of time, occurring when Mozart is in the right emotional state, and when he is undistracted. These ideas are either selected or discarded, and, when retained in memory over some unspecified time, form (again spontaneously?) into completed compositions.

Mozart mentions his prodigious memory, and there is objective evidence of this elsewhere: he was able to transcribe Allegri's *Miserere* from memory after only one hearing, only checking it on the second. His transcription was verified by one of the performers[6]. Mozart identifies his memory as 'perhaps' his 'best gift', maybe a surprising insight for one so consumed by the sound of music. So, arguably, this composer maintained a particular advantage over those endowed with weaker recall: he was capable of conceptualising and then memorising a piece in enough detail that he could 'hear ... the parts ... all at once'. Nevertheless, he writes clearly that imagining is not as good as 'actual hearing'. There are many possible reasons why this might be so, but one plausible account is that there is further elaboration to do as part of the notation process.

---

[5]Edward Holmes, *The Life of Mozart: Including his Correspondence*, Cambridge Library Collection (Cambridge University Press, 2009), pp. 317–8.

[6]Ibid., pp. 66-67.

For many composers (myself included), it doesn't come that easily. The definition of genius, commonly attributed to Thomas Edison, as being five per cent inspiration and ninety-five per cent perspiration is often subverted to describe composition; for example, Steven Stuckey writes:

> You don't make music with ideas, or poetic dreams or wishful thinking. You make it with notes - with technique, with hard work, with Edison's 95 percent perspiration. It is technique that paints paintings, writes poems, builds buildings.[7]

So Stuckey seems to disregard spontaneous, imaginative creativity entirely, not even mentioning the five per cent, though one ought to remark that, in context, this may be for polemic effect. In any case, Mozart's and Stuckey's respective positions serve as proxies for two extremes of a spectrum of opinion: on the one hand, composition is entirely inspiration, and music is born in the mind of the composer (though not necessarily in a single flash of insight); on the other, it is derived from hard graft at the piano or on the page by entirely conscious, reasoned acts of deduction.[8]

Peter Warlock presents a third, different view of composition, placed somewhere between these two poles:

---

[7]Steven Stuckey, `Creating music of geometry and longing', *Cornell University Arts and Sciences Newsletter*, 18/1 (Cornell University, 1996).

[8]An interesting question, which I will expand later, is whether that technique necessarily has to be explicit and conscious, or whether implicit, non-conscious 'technique'—if it can even be called that—is sufficient, as Mozart seems to suggest: for him, musical form apparently 'just happened' without (much) conscious intervention, in the same way as the core ideas.

> If I *had* ideas, I could not write them down without a piano! The sum total of my 'compositions' – (I ought to say 'compilations' for they were all 'discovered' at the piano ...[9]

In Warlock's approach, ideas are 'discovered' at the piano, maybe by improvising in a relatively uncontrolled way and then identifying the 'good', much as Mozart selects some of his imagined ideas and discards others; or maybe by using the piano as a sounding board to work out what the ideas that are imagined actually are in terms of notes. Arguably, *pacet* Stuckey, this is not *only* technique – one might say that the technique visible here is Warlock's 'compilation', and that he is also using his piano to explore the range of imagined possibilities as Mozart's imagination does. So Warlock can be placed on a spectrum somewhere between Mozart and Stuckey. This spectrum allows us to make distinctions between conscious creation in the deliberate planning of a formalist composer, the semi-spontaneous but cooperative and partly planned creation of the jazz improviser in a trio, and the entirely spontaneous whistling in the street of the same people that Schoenberg famously hoped and failed to convince of his 12-note 'tunes'.[10] It is important to note that a non-polar position on this spectrum necessarily entails a *mixture* of explicit technique and implicit imagination: there is not a smooth transition in kind between the two.

---

[9]Barry Smith, (Ed.), *Frederick Delius and Peter Warlock: A Friendship Revealed*, (Oxford University Press, 2000).

[10]Arnold Schoenberg, *Letters*, London: (Faber, 1974). Edited by Erwin Stein. Translated from the original German by Eithne Wilkins and Ernst Kaiser. The word 'tunes' is quoted here because it is Schoenberg's own usage, not because I intend to question its propriety.

So, the hypothesis that I shall put forward in the rest of this chapter proposes two separable, but interacting, cognitive mechanisms involved in composing music, which coexist in such a way as to account for the range of thinking expressed in the spectrum.[11]

Studying Creativity

Before beginning the discussion it will be useful to lay out some theoretical tools. Several general models of creative cognition have been proposed in the past century;[12,13,14,15] each has its own virtue, but for the current purpose, that of Boden is most useful, partly because it can be operationalised mathematically,[16,17] but also because it provides a context in which the other theories, in particular that of Koestler, may be placed.

---

[11] There ought, of course, to be a relationship with improvisation, as variously practised by organists, jazz performers and others, but I shall omit reference to this here, to keep the argument linear.

[12] Graham Wallas, *The Art of Thought*, (New York: Harcourt Brace, 1926).

[13] Arthur Koestler, *The Act of Creation*, (London, UK: Hutchinson, 1976).

[14] Jacob W. Getzels, & Mihaly Csikszentmihalyi, *The Creative Vision: A Longitudinal Study of Problem Finding in Art*, (New York: Wiley, 1976).

[15] Margaret Boden, *The Creative Mind: Myths and Mechanisms*, (London: Weidenfield and Nicholson, 1990).

[16] Geraint A. Wiggins, 'A preliminary framework for description, analysis and comparison of creative systems', *Journal of Knowledge Based Systems*, 19/7 (2006): pp. 449–458.

[17] Geraint A. Wiggins, 'Searching for computational creativity', *New Generation Computing*, 24/73 (2006): pp. 209–222.

Boden's model of creativity revolves around her notion of a *conceptual space* and the exploration of such a space by creative agents – their exact nature is unspecified in the theory: they may be people, computer programs, or other as-yet-unimagined things. The conceptual space is a set of artefacts (in Boden's terms, *concepts*) which are in some quasi-syntactic sense deemed to be acceptable as examples of whatever is being created, so we might take the conceptual space as similar to the set of a certain kind of thing: that which is to be created. Implicitly, the conceptual space may include partially defined artefacts too. *Exploratory creativity* is the process of exploring a given conceptual space, or of selecting an item within it (for example, the range of possible frisbees: different colours, patterns, materials, shapes); *transformational creativity* is the process of changing the rules which delimit the conceptual space (for example, subverting the frisbee to serve as a dinner plate or a hat, or vice versa). Boden[18] also makes an important distinction between mere membership of a conceptual space and the *value* of a member of the space, which is extrinsically defined, but imprecisely. This distinction is easy to see in music: most people can point to pieces of music that they are content to *identify* as such, but which they do not *value* as such; that personal notion of value is easily extended into a collective social construct, also.

An important philosophical point is that the mere existence of the conceptual space does not imply that its contents are known, much as a mathematician's knowledge of the existence of the infinite set of integer numbers does not entail that they have all been written

[18]Margaret Boden, 'Creativity and artificial intelligence', *Artificial Intelligence Journal*, 103 (1998). 347–356.

down. Knowledge of the conceptual space is *intensional*,[19] expressible in terms of properties and constraints rather than by example, and needs to be *extended* to realise concepts from their intensional specification. Given this, exploration of the space becomes something more than mere enumeration of things that are known: it is a little akin to route-finding on a map, in the dark, with a very small torch – one knows the invisible territory exists, but one does not know its form, except by redirecting the torch and looking; on doing this, one can no longer see where one has been.

Bundy[20] and Buchanan[21] join Boden in citing *reflection*, and hence reasoning *about* the conceptual space, rather than *within* it, as a requirement for 'real' or 'significant' creativity (though the definition of such creativity is so far left imprecise). I have shown elsewhere that, in terms of my Creative Systems Framework at least, transformational creativity is precisely exploratory creativity in the conceptual space of conceptual spaces.[22] Lost the thread for 'reflection'…seems to have expired. For completeness, I also mention here that there are other views: Ritchie,[23] for example, presents a completely different account of what is going on in 'transformational' creativity, in which the notion of transformation is

---

[19]Sic. This is the converse of *extensional* and not semantically related to *intentional*.

[20]Alan Bundy, 'What is the difference between real creativity and mere novelty?', *Behavioural and Brain Sciences*, 17/3 (1994): pp. 533–534.

[21]Bruce Buchanan, 'Creativity at the metalevel', *AI Magazine*, 22/3 (2001): pp. 13–28. AAAI-2000 presidential address.

[22] Wiggins, 'A preliminary framework for description, analysis and comparison of creative systems'.

[23]Graeme Ritchie, 'Some empirical criteria for attributing creativity to a computer program', *Minds and Machines*, 17/1 (2007): pp. 67–99.

not so clearly present. Nevertheless, Boden's notion of conceptual space is very helpful to the current discussion.

It is also important not to confuse the dimension of exploratory versus transformational (or object-level versus meta-level[24]) with the dimension of conscious versus non-conscious[25] thought. As a species, humans generally believe they are in much more conscious control than is actually the case. For example, there is evidence that conscious awareness of the intention to speak arises somewhat *after* the commencement of activity associated with generation of linguistic utterances in the brain.[26] Broadly speaking, much of cognition is a good deal less conscious than we tend to think it is, on the basis of introspection, and of course we can only know that which is conscious by definition, so we *would* think that way.

In the computational creativity literature,[27] this introspective bias provides a straw man to sceptics. Proposed mechanisms which might well work at a cognitive level unavailable to conscious introspection, are often derided by human creators because they 'obviously' do not describe what is going on – for example, they say, inspiration, introspected upon, is 'clearly not' the product of systematic enumeration of possibilities. However, the straw man is fireproof, because a majority of cognitive process is not available to introspection. What

---

[24]Wiggins, 'Searching for computational creativity'.

[25]I avoid 'subconscious' to forefend unintended Freudian associations.

[26]Francesca Carota, Andres Posada, Sylvain Harquel, Claude Delpuech, Olivier Bertrand, and Angela Sirigu, 'Neural Dynamics of the Intention to Speak', *Cerebral Cortex*, 20/8 (2010): pp. 1891–7.

[27]See `www.computationalcreativity.net`.

is needed is firstly, to explicitly locate creative cognition with respect to conscious awareness and, secondly, to demystify creativity, accepting that many of the human activities studied in cognitive science are to some extent creative, even if that extent is so small that everyone does it all the time. At this point, we take the controversial step of knocking creativity off its Romantic pedestal.

In this chapter, I use Boden's notion of conceptual space, as characterised above, to capture the set of possible musical compositions. Evidently, this space may be decomposed into smaller spaces capturing different styles, genres or whatever, and I will focus mostly on the conceptual space of tonal melody. Importantly, in my formalisation,[28] the space contains the *empty concept*, a concept with no features at all, which may be thought of as the frame on which all concepts hang, and it also explicitly contains partially defined concepts, which may safely be thought of as pieces in various stages of completeness, though this must be tempered with the acknowledgement that one composer's finished piece may be another's unfinished one. In this way of thinking, a particular compositional path might be described as a point-to-point trajectory from the empty concept (the blank page) to whatever concept corresponds with the final piece. Equally, the trajectory might, star-like, draw together a number of points in the space, each of which begins at a non-empty concept. Here, we are allowing the notion of spontaneous generation of motifs (of whatever kind). In that case, the question must be asked, 'Whence come those initial points?' This is the key question that this chapter aims to answer and the subsequent question begged is, 'How do the points get

[28]Wiggins, 'A preliminary framework for description, analysis and comparison of creative systems'.

joined together?'. I tentatively propose an account for that also. First, however, we must ask whence for the individual comes the conceptual space?

Learning and Creativity

To find the source of the conceptual space, we must examine the human capacity to learn. To deny this relation would be to deny the evidence observed by every parent in history: we are not born with a full understanding of the world around us, but we must learn it. The vast majority of learning, however, is not done explicitly at school or from books, but implicitly, from direct experience of the world itself. In the musical context, even capacities such as the perception of relative pitch (as distinct from the trained ability to name the interval classes thus formed) are implicitly learned,[29] and are learned with surprising efficiency, given exposure to the necessary stimuli. Musical style is self-evidently learned too: an individual enculturated in Africa, for example, has a different internal model of musical style from an individual extensively and exclusively exposed to European music[30]. In the complete absence of evidence for genetic encoding of musical style, and given the substantial and increasing body of evidence for implicit musical learning, it is the latter hypothesis that is more convincing. The existential, evolutionary value of the voracious

---

[29]Jenny Saffran and Gregory Griepentrog, 'Absolute pitch in infant auditory learning: Evidence for developmental reorganization', *Developmental Psychology*, 37/1 (2001): pp. 74–85.

[30]Petri Toiviainen, and Tuomas Eerola, `Where is the beat?: Comparison of Finnish and South-African listeners' in Reinhard Kopiez, Andreas Lehmann, Irving Wolther and Christian Wolf, (Eds.), *Proceedings of the 5th Triennial ESCOM Conference* (ESCOM, 2005).

human capacity to learn[31] is based in its affordance of expectation: massively enhanced ability to manage the world,[32] which is strictly necessary for such a frail organism.

It is a property of the human mind to generalise from examples: indeed, the tendency is so strong as sometimes to lead us astray in quite simple logic, contributing to major social problems such as racism. At the non-conscious perceptual level, generalisation allows us to capture the essential properties of events and objects in the world around us, and serves to protect us from threats which are similar to, but not the same as, threats previously encountered. In the case of music, as we are exposed to more and more examples, the more we tend to generalise the style, the more we learn about it, and the more we develop efficient ways of hearing and remembering it. For example, experienced listeners to music of the Classical period can develop a very strong sense of tonal structure, and are able to hear movement around the various tonal functions; this is not the same as being taught the music-theoretic tonal functions and knowing about them. The two kinds of knowledge are quite independent. Learning, for example, what a dominant sounds like is cognitively useful in the face of large quantities of tonal data because it helps promote cognitively efficient classification and recall of that sound data, even if the learner-perceiver does not know that this is called 'dominant' by music theorists. It follows that someone never previously

---

[31] Irving Biederman and Edward Vessel, `Perceptual pleasure and the brain', *American Scientist*, 94 (2006): pp. 247–53.

[32] Marcus T. Pearce and Geraint A. Wiggins, 'Auditory expectation: The information dynamics of music perception and cognition', *Topics in Cognitive Science* (Wiley, 2011, in press).

exposed to a particular style is probably cognitively incapable of hearing the finer-grained aspects of that style that an experienced listener enjoys.

In context of Boden's theory, this implicitly-learned, generalised perceptual mental model of music is one candidate to supply the conceptual space. The other candidate is a theoretically-acquired, explicitly-learned model of music, as taught in music theory classes across the Western world. On the compositional level, these two possibilities correspond with the two ends of my compositional spectrum, introduced earlier; the former is at the Mozart-like, spontaneous end, while the latter is Stuckey-like, relying on technique and knowledge *about* music, rather than implicit musical imagination. As might be expected, a mixture of the two can produce a Warlock-amalgam of behaviours, where explicit and implicit knowledge interact. However, when we are learning to *hear* music, there is only the perceptually-learned model: the music-theoretic account is meaningless, in a literal sense, if one has not learned the necessary cognitive representations for the style in question. It is an interesting point to note that this meaninglessness does not prevent the successful application of theoretical rules, at least at a simple level. This is the same property that allows logical calculi to propose solutions to problems in reasoning: their syntax directly encodes their semantics.

A model of musical learning that can perhaps create

One way to provide evidence for the hypothesis developed here as an account of human creativity is to build computer models of it, and demonstrate that behaviour of such models predicts that of human creators. To do so, we must begin with a model of learning – here, I use that developed/proposed/formalised by Pearce[33], a complex and detailed model of auditory sequence learning, embodied in a computer program. The detail of the model and the program that embodies it are not relevant to the current argument, except in the following. The model has no programmed rules about musical style, though it has the simulated capacity to perceive[34] various musical constructs, such as scale degree and key note. The model is exposed to a large body of tonal melodies from which it learns, merely by counting the number of occurrences of each kind of event in sequence, in context of what preceded it, using various different representations (for example, absolute pitch, scale degree, note duration) simultaneously to do so. The different representations predict separately, but their predictions are combined into one for each note, using Shannon's mathematical information theory.[35,36] Finally, and crucially, the model is able to generalise,

---

[33]Marcus T. Pearce, *The Construction and Evaluation of Statistical Models of Melodic Structure in Music Perception and Composition*, PhD thesis, Department of Computing, City University, London, UK, (2005).

[34]Avoiding anthropomorphism in this kind of discussion, as would be ideal, entails awkward and continual circumlocution, or at best an excess of quotation marks. It clearest simply to use the anthropomorphic terminology while reminding ourselves that all of this quasi-human traits are *simulated*.

[35]Claude Shannon, 'A mathematical theory of communication', *Bell System Technical Journal*, 27 (1948): pp. 379–423, 623–56.

[36]Marcus T. Pearce, Darrell Conklin and Geraint A. Wiggins, 'Methods for combining statistical models of music', in Uffe Kock Wiil (Ed.), *Computer Music Modelling and Retrieval*. (Heidelberg, Germany: Springer Verlag, 2005), pp. 295–312.

to accommodate events that it has not previously encountered;[37] again the detail of how is not relevant here. Once learning is complete, the model is able to predict the expectations of listeners enculturated into Western music to a surprising degree of accuracy: in statistical terms, it accounts for up to 81 per cent of the variance in human responses[38] (of course, not all listeners respond identically, so there is no single correct answer). The model is also able to predict *segmentation* of musical melody from its learned data alone., the points at which phrase boundaries are perceived by listeners, to a degree comparable with explicitly music-theoretical approaches.[39]

Because it is capable of successful application to these tasks, I hypothesise that the model may serve as a simulated conceptual space for the melodies it learns, and, further, that the simulation is of a human conceptual space, not an arbitrary computational one. This claim is underpinned by the empirical evidence cited here for the behaviour of the model as a model of perception, and by the well-supported hypothesis that a key feature of human conceptual spaces is their close relationship with perception.[40]

---

[37]Marcus T. Pearce and Geraint A. Wiggins, 'Expectation in melody: The influence of context and learning', *Music Perception*, 23/5 (2006): pp. 377–405.

[38]Ibid.

[39]Marcus T. Pearce, Daniel Müllensiefen and Geraint A. Wiggins, 'The role of expectation and probabilistic learning in auditory boundary perception: A model comparison', *Perception*, 9 (2010): pp. 1367–1391.

[40]Peter Gärdenfors, *Conceptual Spaces: the geometry of thought* (Cambridge, MA: MIT Press, 2000).

This model, however, is rather abstract; it is not a direct model of the neural behaviour of the brain (though it does seem to have certain neural correlates[41]). This, however, does not undermine its status as a model of cognitive function: the proof is in its demonstrable ability to predict human behaviour. Because it does so by counting observed occurrences it is, in essence, a statistical model. This means that, in principle, standard methods from statistics can be used to sample from the model's memory, to produce complete melodies in the broad style that the model has learned.This has been shown to work, to a minimal level of acceptability, for melodies in the style of those harmonised as chorales by Johann Sebastian Bach.[42] The melodies are rarely good, but they are recognisable as melodies, when rigorously evaluated by independent observers. The one selected as best by our observers is reproduced in Figure **Error! Reference source not found.**1.

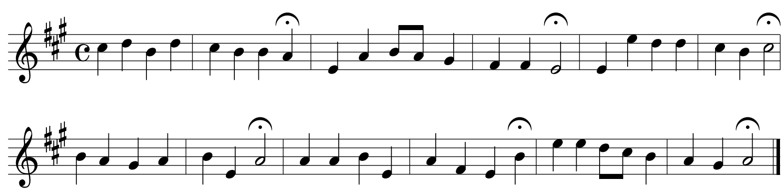

Figure 1: A statistically-generated chorale melody, deemed acceptable by human listeners. The rhythm is taken from *Jesu, meiner Seelen Wonne* (BWV 359); pitches are generated by our statistical model without human intervention.

[41]Marcus T. Pearce, Maria Herrojo Ruiz, Selina Kapasi, Geraint A. Wiggins and Joydeep Bhattacharya, 'Unsupervised statistical learning underpins computational, behavioural and neural manifestations of musical expectation', *NeuroImage*, 50/1 (2010): pp. 303–314.

[42]Marcus T. Pearce and Geraint A. Wiggins, 'Evaluating cognitive models of musical composition', in Amílcar Cardoso and Geraint A. Wiggins (Eds.), *Proceedings of the 4th International Joint Workshop on Computational Creativity* (2007), pp. 73–80.

Whorley has developed a more advanced statistical system[43,44] that is capable of harmonising hymn tunes to a reasonably high musical level, demonstrating the generality of the ideas beyond melody alone; I give an example in Figure **Error! Reference source not found.**2. Substantial work remains to be done before strong claims can be made, however.

Figure 2. A harmonisation by Whorley's statistical harmonisation system, again without human intervention. The tune is a French church melody, from *Chants*

[43]Raymond Whorley, Marcus T. Pearce, and Geraint A. Wiggins, 'Computational modelling of the cognition of harmonic movement', in *Proceedings of the 10th International Conference on Music Perception and Cognition*, Sapporo, Japan (2008).

[44]Raymond Whorley, Geraint A. Wiggins, Christophe Rhodes and Marcus T. Pearce, 'Development of techniques for the computational modelling of harmony', in Ventura et al. (Eds.), *Proceedings of the First International Conference on Computational Creativity* (2010).

***Ordinaires de l'Office Divin* (Paris, 1881); it is reprinted as Hymn No. 33, *Grafton*, in the 1993 edition of the English Hymnal.**

Returning briefly to my compositional spectrum, the sampling approach described above is an extreme case of Mozart-myth, in that it supposes a complete melody appearing in one flash of sampled inspiration, equivalent to picking a completed piece out of the conceptual space as one point, fully formed, as Minerva from the forehead of Jove. Therefore, I do not propose it as a representative model of human composition. A closer simulation of Mozart's self-reported approach would be the statistical generation of melodic fragments, or motifs, which are glued together by subsequent traversal of the conceptual space – the star-shaped creative trajectory mentioned in an earlier section. But how can statistical sampling (over whole pieces or fragments) be justified as a cognitive model of creativity?

I have already mentioned the human property of expectation, casting it in general as a device for managing the world. But what is the relation between statistical memory of musical melody and expectation? The answer is quite simple: one non-consciously predicts what happens next on the basis of what one has experienced in the immediate past. This is a very efficient way of managing information in the world, because relevant memories can be accessed in advance, priming us to be ready for what is next. What is more, appropriate cognitive processing power can be applied: something that is expected needs less processing – less attention – than something which is not, because we already know what it is. So, simply put, we expect things more in a given context if we have heard something similar in a similar context before; we are non-consciously, continually guessing what comes next. This simple idea (which is more complicated to implement) seems to underlie several aspects of speech processing as well – indeed, the model proposed here is capable of segmenting speech into morphemes, using the same method it uses for phrase segmentation.[45] Most importantly, predicting what comes next helps us communicate more efficiently. So it is established that, when listening to music and speech, there is a cognitive process which continually predicts what is coming next. It is entirely reasonable to propose that this process equally capable of responding to internal (imagined) musical phenomena as external (heard) ones and so it is not hard to imagine a situation where any sound, real or remembered, might trigger the expectation mechanism, suggesting a continuation into a sequence. Given the position that learning is an essentially statistical process, to do with correlations of co-occurrence in observed events, the expectation mechanism can be thought of as statistical sampling: the generation of instances from a statistical model. Crucially, the involvement of generalisation means that it is possible to sample instances that the model has not specifically observed.

In this more incremental view, we might see small sections of music – motifs – appearing, note by note, rather than as a whole, pre-formed melody. As a result we can generate small units, and consider their likelihood in terms of the extant model: the non-conscious correlate of this latter activity being the ability to 'see' (that is, without conscious consideration) how a motif might fit into one or another context. Following the *hedonic*

[45] Geraint A. Wiggins, ' "I let the music speak": cross-domain application of a cognitive model of musical learning', in Patrick Rebuschat and John Williams (Eds.), *Statistical Learning and Language Acquisition* (Amsterdam, NL: Mouton De Gruyter, 2011, in press).

curve of Wilhelm Wundt[46,47] (see Figure **Error! Reference source not found.**3), we find that very likely units are dull, while very unlikely ones are difficult to relate to the conceptual space – exactly as we might expect from the corresponding probabilities. There is a balance to be struck between novelty and stylistic conformity. Each of these units corresponds with a point, a non-empty concept, in the conceptual space, just as Mozart's initial ideas gave him starting points for musical imagination. This, then, is a hypothetical account of how inspiration, at the level of motivic ideas, might happen. The next step in my argument is to propose why it happens in the cognitive context.

The chattering crowd of mind

For decades, in Artificial Intelligence, the notion of collective, agent-based models of mind have been current – for example, Marvin Minsky's famous Society of Mind[48]. Here the massively parallel nature of cognitive processing is captured in systems which consist of interacting processes, whose combined emergent properties are then complex and unpredictable. Within this broad category, Bernard Baars has proposed Global Workspace Theory,[50] a convincing theory of consciousness based on information production by large

---

[46]Wilhelm Wundt, *Grundzüge der physiologischen Psychologie*. (W. Engelman, 1874). Translated into English as Wilhelm Wundt and Edward Titchener, *Principles of physiological psychology*, volume 1 of *Principles of Physiological Psychology* (Sonnenschein, 1904).

[47]Elizabeth Margulis and Andrew Beatty, 'Musical style, psychoaesthetics, and prospects for entropy as an analytic tool', *Computer Music Journal*, 32/4 (2008): pp. 64–78.

[48] Marvin Minsky, *The society of mind* (New York: Simon and Schuster, Inc., 1985)

[50]Bernard Baars, *A cognitive theory of consciousness*. (Cambridge University Press, 1988).

numbers of cognitive processes, which may operate in various degrees of synchrony, and high synchrony is associated with conscious awareness. Murray Shanahan gives neurophysiological underpinning for the theory.[53] It is impossible to give a complete account of these wide-ranging, subtle and elegant ideas in the space available here, so I summarise instead by analogy. Human cognition may be conceptualised as a crowd of book-makers,[54] each of which continually shouts the odds of informational tips from data provided by sensory mechanisms and memory. Some of the tips make it as far as the Global Workspace, where they can be heard by everyone, but some of them are lost on the way, in the constant babble of shouting book-makers. When a tip from one bookie appears in the Global Workspace, it becomes accessible, not only to all the others, but also to consciousness, by a process which remains somewhat obscure – but that is not the focus of the current argument.

To match the theory against Mozart's reported creative experience, a point-by-point summary is useful. To deal with the world, we constantly predict from models built statistically from experience. Prediction is continual, multiple and in parallel. Predicted items are selected, and made available to consciousness, so there is a notion of competition between predictors. When a prediction is selected, the new information becomes available to all predictors. Now, let us compare this sequence of events with Mozart's report. When

---

[53]Murray Shanahan, *Embodiment and the inner life: Cognition and Consciousness in the space of possible minds*. (Oxford Univeristy Press, 2010).

[54]That is to say, not publishers, but people who give the odds, take bets and offer 'tips'—not gratuities, but suggestions for promising bets—on horse races.

he is ‘completely himself’, when nothing is distracting him, and his mind is open, so there is relatively little to compete with musical predictors, ideas ‘come’, spontaneously (he ‘cannot force them’). The ideas cannot be ‘forced’, so the essence of the process is non-conscious. Some of the ideas formed do not ‘please’ him, and are deliberately discarded, so we may hypothesise that whatever non-conscious ‘selection’ is applied at this stage of Mozart’s compositional process is not a complete determinant of musical value, just like our experimental melody creators, mentioned above.

Having been selected, first by the mysterious Global Workspace selection mechanism, and then by Mozart’s own idiosyncratic hedonic assessment (the details of which are inscrutable, since it passed away with its owner), the chosen ideas are consciously memorised and therefore available to all predictors in the Global Workspace. Given the overwhelming tendency to predict, what would be more natural than to predict musical structures that include these multiple smaller units, appropriately connected together? In this way, each prediction can step towards a completed composition, guided but not restricted by the hypothetical composer’s generalised statistical model of style, with the composer selecting or manipulating at each cycle: ‘my subject enlarges itself, becomes methodized and defined.’ One can also speculate that a composer with Mozart’s capacious and punctilious memory might be able to remember the process as well as the outcome; and the sequence of events proposed here does seem to match his introspective description ,though, of course, there may be other candidates, and the question of which really is correct cannot be answered.

There remain two outstanding aspects of the current matter for which I have not proposed hypothetical solutions: first, the conscious selection of potentially usable motifs on the basis of idiosyncratic quality, and, second, the non-conscious restriction of the predictors' output.

We can seek a solution to the first of these in the relatively new field of neuroaesthetics, where, for example, Biederman & Vessel[55] have given convincing proposals and evidence as to how somatic hedonic responses may be derived from the process of learning about perceived objects. This work is in its infancy, and is an area where great scientific contributions to the understanding of humankind are to be made. We can, also, in more conscious contexts (for example, the deliberate construction of rock 'anthems'), imagine a rule-based selector, working on the basis of stylistic similarity to analytically-identified key features of other music of the target kind.

A candidate solution for the second question (of how non-conscious selection is applied) can be given directly, in terms of the operation of the statistics of the predictors, in terms of their memory models, but also in terms of the collective crowd. There are two factors involved here. The first is the probability of the item being predicted and, relatedly, the amount of information it carries. As one is sampling continually, and more probable items are likely to be selected more frequently, so we might expect the Workspace to be cluttered with banal rubbish. However, Wundt[56] tells us that likely items are less interesting than

[55]Biederman and Vessel, 'Perceptual pleasure and the brain' and later work.

[56]Wundt, *Grundzüge der physiologischen Psychologie*.

unlikely ones; in terms of Shannon's[57] information theory, they do not carry much information. So part of my candidate solution is that there needs to be a certain amount of information in a prediction before it is allowed through into the global workspace. This is philosophically reasonable: there would be little value in conscious awareness of prediction if it were not to isolate the unexpected; valued cognitive resources would be wasted. There is empirical evidence for this in perception: we[58] have shown that there is indeed increased global synchrony (which, I noted above, corresponds in Global Workspace Theory with conscious awareness) in beta-band neural activity, in response to improbable melodic notes, as compared with probable ones. I propose that the same mechanism operates on produced items as on perceived ones.[59] Thus, information-heavy predictions are preferred over information-light ones.

---

[57]Shannon, 'A mathematical theory of communication'.

[58]Pearce, Herrojo Ruiz, Kapasi, Wiggins and Bhattacharya, 'Unsupervised statistical learning underpins computational, behavioural and neural manifestations of musical expectation'.

[59]Here, as elsewhere, I appeal to Occam's razor, a scientific precept due to William of Ockham: a simple theory is better than a complicated one. In the current case, one mechanism applicable in two instances is simpler than two different mechanisms, one for each instance.

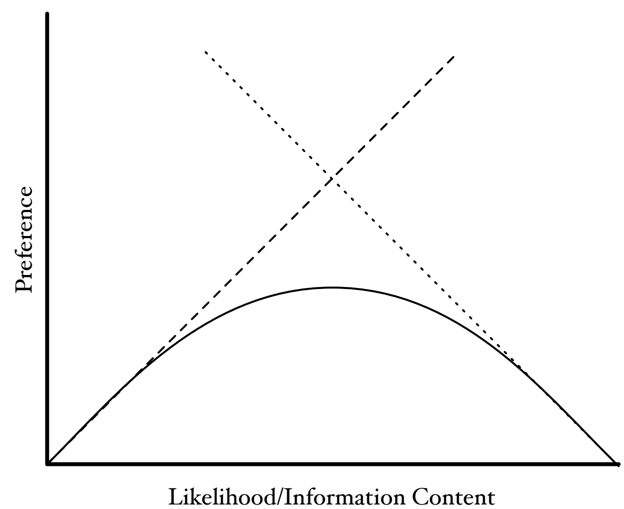

Figure 3. The Wundt curve (solid) is formed by the multiplication of two linear functions: the likelihood of a generated item (dashed) and the number of generators likely to agree on an item, according to its likelihood (dotted).

The second part of this second solution emerges straightforwardly from the statistical dynamics of the system. According to Shanahan,[60] the synchrony that corresponds with availability to consciousness can be thought of as multiple predictors producing a particular solution simultaneously. At any given time, and given a large number of randomly sampling predictors, likely predictions will be more common in the population; unlikely ones, which contain more information, less common, and outrageously obscure ones very rare indeed. So it is less likely that sufficient synchrony will be achieved for conscious awareness, by the less likely outcomes, simply because it is less likely that they will be selected. This gives us two opposing trends, which are linear functions of likelihood, one positive and one negative, as illustrated in Figure **Error! Reference source not found.**.3. The multiplicative

[60]Shanahan, *Embodiment and the inner life: Cognition and Consciousness in the space of possible minds*.

combination of these two, motivated by the usual combination of independent likelihoods in probability theory, gives (one version of) the Wundt curve, as shown in the figure. This explanation is powerful, because it needs no mechanism other than that already proposed above and elsewhere for other purposes.

At this stage, I have proposed a hypothetical mechanistic cognitive framework in which inspiration can take place.

Technique, structure and how to 'see the truth'

I now return to the question of technique, and its opposition to inspiration. But the first question to ask is: 'Are these two really in opposition?' To help consider this question, imagine, on one hand, a music freshman laboriously harmonising a melody in the style of Bach according to taught rules on paper, and, on the other, an expert organist, harmonising the same melody in the same style, live at the keyboard, during a Sunday service. While the organist certainly has knowledge of the rules that the student is applying, she does not need to think very hard, if at all, about them: she simply *feels* what comes next, to the extent that her fingers almost seem to work independently from conscious intervention. Karmiloff-Smith[61] identifies three cognitive stages of learning: a data-driven acquisition phase, where representations are independently stored and used, with no internal representational change; a middle phase of internal representational change, during internalisation of knowledge; and finally, reconcilation of the new internal representations with the external knowledge, which includes reflection. Our student is at the first, following rules, while our organist has

[61] Annette Karmiloff-Smith, *Beyond Modularity: A Developmental Perspective on Cognitive Science*. (Cambridge, MA: The MIT Press, 1995).

reached the third, where knowledge is encoded so deeply in the non-conscious system that the effect of using it is encoded in common parlance. In other words, she has the music ‘under her fingers.’ Mozart, it appears, was at this level in terms of imagining music. Warlock, on the other hand, maybe used his implicit deep knowledge of piano technique to guide his ‘discovery’ of new musical ideas, but without Mozart’s level of imagination of music. Again, it regrettably too late to test this proposal.

Karmiloff-Smith’s three stages of learning constitute a comparator against which a given individual’s capacity at a given task can be measured, broadly summarised as the range between completely explicit reasoning (conscious, rule-based) and completely implicit reasoning (non-conscious, intuitive). For the music freshman with the necessary inclination, diligence and core ability, there is a trajectory to the level of the organist, moving along this spectrum – and that is the original point of the theory.

The Stuckey-end of my compositional spectrum seems at first sight to correspond to some degree with Karmiloff-Smith’s first level (conscious application of rules), but this is misleading – it is altogether more complicated. Compositional technique (as distinct from harmonisation exercises where rules are predefined) is not merely about the application of given rules, it also entails intuition or design of the rules to be applied, and their effective application. These actions are transformational, in Boden’s sense, so may be considered ‘more’ creative than mere application of taught rules.

Computational rule-based systems have been applied extensively to compositional tasks in the past,[62,63] sometimes explicitly from the point of view of *search control*,[64,65] which may

[62]David Cope *Computers and Musical Style*. (Oxford University Press, 1991).

be thought of as technique: 'Which rule(s) should I apply at each point?' The problem, though, with such methods, is that one always needs to start from somewhere: even given an initial empty concept, a first step has to be taken. In harmonisation examples, there is the given melody and all the implicit framing-information it brings with it (key, tonality, and so on). In David Cope's EMI, it seems that structure from human-composed music supplies the basis.[66] It is hard to see indeed how a rule-based system can start from a blank page, unless it is by fiat from a programmer (so the seed of creativity is coming from a human, not the program, and therefore the program is not embodying a complete theory of creativity) or by random selection, as practiced (usually with a non-uniform distribution, derived from prior musical knowledge) in evolutionary music generation.[67]

---

[63]George Papadopoulos and Geraint A. Wiggins, 'AI methods for algorithmic composition: A survey, a critical view and future prospects', in *Proceedings of the AISB'99 Symposium on Musical Creativity* (Brighton, UK: SSAISB, 1999), pp. 110–117.

[64]Kemal Ebcioğlu, 'An expert system for harmonizing four-part chorales', *Computer Music Journal*, 12/3 (1988): pp. 43–51.

[65]Somnuk Phon-Amnuaisuk, Alan Smaill and Geraint A. Wiggins, 'Chorale harmonization: A view from a search control perspective', *Journal of New Music Research*, 35/4 (2006): pp. 279–305.

[66]David Cope *Virtual Music: Computer Synthesis of Musical Style* (Cambridge, MA: MIT Press, 2004).

[67]Geraint A. Wiggins, George Papadopoulos, Somnuk Phon-Amnuaisuk and Andrew Tuson, 'Evolutionary methods for musical composition', *International Journal of Computing Anticipatory Systems* (1999).

Towards an objective account of creative inspiration

This idea brings us full circle, since biased random sampling, followed by selection, is precisely the cognitive mechanism proposed above. I have cited some preliminary empirical evidence for the mechanisms proposed, too, suggesting that one can have at least some confidence in the proposal – there is no pretence at this stage of more than hypothetical claims.

Given such a hypothesis, one can in principle begin to build a veridical simulation of the creative mechanisms proposed here, and music is an ideal domain in which to work, being untrammelled by reasoning about the real-world that makes linguistic and non-abstract visual art difficult for creative computer systems.

Most importantly, from these precepts, we can at last, begin to give an account of human creative inspiration in music, in terms of conscious and non-conscious processes, some of which are so banal as to go unnoticed every day, and some of which can produce humankind's greatest art.

References

Baars, Bernard J. *A cognitive theory of consciousness* (Cambridge, UK: Cambridge University Press, 1988).

Biederman, Irving and Edward A. Vessel, 'Perceptual pleasure and the brain', *American Scientist*, 94 (2006): pp. 247–53.

Boden, Margaret A., *The Creative Mind: Myths and Mechanisms* (London: Weidenfield and Nicholson, 1990).

Boden, Margaret A., 'Creativity and artificial intelligence', *Artificial Intelligence Journal*, 103 (1998): pp. 347–356.

Buchanan, Bruce G., 'Creativity at the metalevel', *AI Magazine*, 22/3 (2001): pp. 13–28. AAAI-2000 presidential address.

Bundy, Alan, 'What is the difference between real creativity and mere novelty?', *Behavioural and Brain Sciences*, 17/3 (1994): pp. 533–534.

Carota, Francesca, Andres Posada, Sylvain Harquel, Claude Delpuech, Olivier Bertrand, and Angela Sirigu, 'Neural Dynamics of the Intention to Speak', *Cerebral Cortex*, 20/8 (2010): pp. 1891–7.

Cope, David, *Computers and Musical Style*, (Oxford: Oxford University Press, 1991).

Cope, David, *Virtual Music: Computer Synthesis of Musical Style*, (Cambridge, MA: The MIT Press, 2004).

Ebcioğlu, Kemal, ‘An expert system for harmonizing four-part chorales’, *Computer Music Journal*, 12/3 (1988): pp. 43–51.

Gärdenfors, Peter, *Conceptual Spaces: the geometry of thought*, (Cambridge, MA: The MIT Press, 2000).

Getzels, Jacob and Mihalyi Csikszentmihalyi, *The Creative Vision: A Longitudinal Study of Problem Finding in Art*, (New York: Wiley, 1976).

Holmes, Edward, *The Life of Mozart: Including his Correspondence*, Cambridge Library Collection, (Cambridge, UK: Cambridge University Press, 2009).

Karmiloff-Smith, Annette, *Beyond Modularity: A Developmental Perspective on Cognitive Science*, (Cambridge, MA: The MIT Press, 1995).

Koestler, Arthur, *The Act of Creation*, (London, UK: Hutchinson, 1976).

Margulis, Elizabeth H. and Andrew P. Beatty, ‘Musical style, psychoaesthetics, and prospects for entropy as an analytic tool’, *Computer Music Journal*, 32/4 (2008): pp. 64–78.

Papadopoulos, George and Geraint A. Wiggins, ‘AI methods for algorithmic composition: A survey, a critical view and future prospects’, in *Proceedings of the AISB’99 Symposium on Musical Creativity*, (Brighton, UK: SSAISB, 1999), pp. 110–117.

Pearce, Marcus T. *The Construction and Evaluation of Statistical Models of Melodic Structure in Music Perception and Composition*. PhD thesis, Department of Computing, City University, London, UK (2005).

Pearce, Marcus T., Darrell Conklin, and Geraint A. Wiggins, ‘Methods for combining statistical models of music’, in Uffe Kock Wiil (Ed.), *Computer Music Modelling and Retrieval,* (Heidelberg, Germany: Springer Verlag, 2005), pp. 295–312.

Pearce, Marcus T., Maria Herrojo Ruiz, Selina Kapasi, Geraint A. Wiggins and Joydeep Bhattacharya, ‘Unsupervised statistical learning underpins computational, behavioural and neural manifestations of musical expectation’, *NeuroImage*, 50/1 (2010): pp. 303–314.

Pearce, Marcus T., Daniel Müllensiefen, and Geraint A. Wiggins, ‘The role of expectation and probabilistic learning in auditory boundary perception: A model comparison’, *Perception*, 9 (2010): pp. 1367–1391.

Pearce, Marcus T. and Geraint A. Wiggins, ‘Expectation in melody: The influence of context and learning’, *Music Perception*, 23/5 (2006): pp. 377–405.

Pearce, Marcus T. and Geraint A. Wiggins, ‘Evaluating cognitive models of musical composition’, in Amìlcar Cardoso and Geraint A. Wiggins (Eds.), *Proceedings of the 4th International Joint Workshop on Computational Creativity* (2007), pp. 73–80.

Pearce, Marcus T. and Geraint A. Wiggins, ‘Auditory expectation: The information dynamics of music perception and cognition’, *Topics in Cognitive Science*, (2011, in press).

Phon-Amnuaisuk, Somnuk, Alan Smaill and Geraint A. Wiggins ‘Chorale harmonization: A view from a search control perspective’, *Journal of New Music Research*, 35/4 (2006): pp. 279–305.

Ritchie, Graeme, ‘Some empirical criteria for attributing creativity to a computer program’, *Minds and Machines*, 17/1 (2007): pp. 67–99.

Saffran, Jenny R. and Gregory J. Griepentrog, ‘Absolute pitch in infant auditory learning: Evidence for developmental reorganization’, *Developmental Psychology*, 37/1 (2001): pp. 74–85.

Schoenberg, Arnold, *Letters*. Edited by Erwin Stein. (London: Faber, 1974). Translated from the original German by Eithne Wilkins and Ernst Kaiser.

Shanahan, Murray, *Embodiment and the inner life: Cognition and Consciousness in the space of possible minds*. (Oxford: Oxford University Press, 2010).

Shannon, Claude, A mathematical theory of communication, *Bell System Technical Journal*, 27 (1948): pp 379–423, 623–56.

Smith, Barry (Ed.). *Frederick Delius and Peter Warlock: A Friendship Revealed*. (Oxford: Oxford University Press, 2000).

Stuckey, Steven, ‘Creating music of geometry and longing’, *Cornell University Arts and Sciences Newsletter*, 18/1 (1996).

Toiviainen, Petri and Tuomas Eerola, ‘Where is the beat? : Comparison of Finnish and South-African listeners’, in Reinhard Kopiez, Andreas C. Lehmann, Irving Wolther and Christian Wolf, (Eds.), *Proceedings of the 5th Triennial ESCOM Conference*, (ESCOM, 2005).

Wallas, Graham, *The Art of Thought*, (New York: Harcourt Brace, 1926).

Whorley, Raymond, Geraint A. Wiggins, Christophe Rhodes, Marcus T. Pearce, ‘Development of techniques for the computational modelling of harmony’, in Ventura et al. (Eds.), *Proceedings of the First International Conference on Computational Creativity* (2010).

Whorley, Raymond, Marcus T. Pearce, and Geraint A. Wiggins, 'Computational modelling of the cognition of harmonic movement', in *Proceedings of the 10th International Conference on Music Perception and Cognition* (2008).

Wiggins, Geraint A., 'A preliminary framework for description, analysis and comparison of creative systems', *Journal of Knowledge Based Systems*, 19/7 (2006): pp. 449–458.

Wiggins, Geraint A., 'Searching for computational creativity', *New Generation Computing*, 24/3 (2006): pp. 209–222.

Wiggins, Geraint A., 'Semantic Gap?? Schemantic Schmap!! Methodological considerations in the scientific study of music', in *Proceedings of 11th IEEE International Symposium on Multimedia* (2009), pp. 477–482.

Wiggins, Geraint A., ' "I let the music speak": cross-domain application of a cognitive model of musical learning', in Patrick Rebuschat and John Williams (Eds.), *Statistical Learning and Language Acquisition* (Amsterdam, NL: Mouton De Gruyter, 2011, in press).

Wiggins, Geraint A., Daniel Müllensiefen and Marcus T. Pearce, 'On the non-existence of music: Why music theory is a figment of the imagination', *Musicae Scientiae*, Discussion Forum 5 (2010): pp. 231-255.

Wiggins, Geraint A., George Papadopoulos, Somnuk Phon-Amnuaisuk and Andrew Tuson, 'Evolutionary methods for musical composition', *International Journal of Computing Anticipatory Systems* (1999).

Wroe, Nick, ‘ A life in music: Richard Rodney Bennett’, *The Guardian* (22nd July, 2011).

Wundt, Wilhelm and Edward Titchener, *Principles of physiological psychology*, vol. 1. (Sonnenschein, 1904).

Wundt, Wilhelm *Grundzüge der physiologischen Psychologie*. (W. Engelman, 1874).