Computer Models of Musical Creativity: A Review of Computer Models of Musical Creativity by David Cope

Correspondence:
Geraint A. Wiggins,
Department of Computing,
Goldsmiths College,
University of London,
New Cross, London
SE14 6NW, UK.
E-mail: g.wiggins@mac.com

Geraint A. Wiggins
Centre for Cognition, Computation and Culture, Goldsmiths College, University of London, UK

1 Introduction

Pseudoscience (as defined by Wikipedia on 4th July 2007) is

> any body of knowledge, methodology, belief, or practice that claims to be scientific or is made to appear scientific, but does not adhere to the basic requirements of the scientific method.

(http://en.wikipedia.org/wiki/Pseudoscience) Examples from the same article are 'astrology, quackery, the occult, and superstition', and some identifying features are (paraphrased):

(1) Use of vague, exaggerated or untestable claims—including: (a) assertion of imprecise claims and (b) use of obscurantist language.
(2) Over-reliance on confirmation rather than refutation—including: (a) assertion of claims that cannot be falsified in the event that they are incorrect; (b) over-reliance on testimonials and anecdotes; (c) selective use of experimental evidence.
(3) Lack of openness to testing by other experts—including: failure to provide adequate information for other researchers to reproduce the claimed results.
(4) Lack of progress—including: failure to progress towards additional evidence of its claims.
(5) Personalisation of issues—including: (a) tight social groups and granfalloons;[1] (b) attacking the motives or character of anyone who questions the claims.

When one sets out to study or replicate human or machine creative behaviour in anything like a scientific way, one must be very careful, because creativity is shrouded in both mystery and mystique. Psychologically and philosophically, creativity is every bit as difficult to identify as consciousness. Proponents of artificial intelligence often shy away from computational creativity as the inscrutable, unachievable thing, much in the same way that their opponents refuse to accept the idea of computational intelligence. Aesthetically, creativity is extremely hard to define: it is relative to both temporal and social context, it can apply at any number of different but interrelated levels (e.g. perceptual, narrative, theoretical). In general, we can recognise something as creative when we see it, and even explain individual cases, but we cannot generally give rules as to what is really a product of creative behaviour. Finally, of course, to many people, creativity is a precious, personal thing, almost like a religious construct, which is not open to question or even scrutiny.

David Cope's latest book, *Computer Models of Musical Creativity* (2005), sets out to study musical creativity, which, he argues, is different from other

doi:10.1093/llc/fqm025 Advance Access Published on 5 November 2007

kinds. This argument rests mostly on the claim that music is a different kind of thing from other things in the production of which creativity might be exhibited. For example, the argument runs, musical creativity is different from creativity in language (the '*sine qua non*' of language use, according to Plotkin, 1998) because language has referential semantics and music (mostly) does not. This latter debate has been had elsewhere (e.g. Wiggins, 1998), but it is not at all clear that the inference connecting it to the nature of musical creativity is a valid one. Cope merely asserts the inference, and does not defend it; Boden (1990, 1998), on the other hand, gives a model of creativity which is general, and Wiggins (2001) shows how it can apply to music. While this fundamental lack of philosophical basis does not ultimately damage the outcome, as will become clear below, it does weaken one's faith in the argument.

The book is something of a chimera. Cope apparently sets out to be scientific. The cover notes say that the book is about 'experimental' work. The title refers to 'computer models'; its contribution is described with words like 'study'; and the author says he sets out to 'prove' things (see *Preface*)—all of which are generally thought of as parts of a scientific approach. It includes descriptions (albeit mostly imprecise ones) of mathematical processes, it lists formulae, and gives very complicated diagrams of so-called 'association networks', more on which below. I would argue that everything about the book suggests that the work it outlines is intended to be viewed as something scientific; but later (on page 356 of 375 of main text) we are told that '*what matters most is the music*' (original emphasis). This leaves me in a difficult position in reviewing the book. Do I review what it seems to be (a scientific study of computer models of creativity), or what the author (says he) means it to be (supporting material for his music?)? Literary critics tell us that the latter is the *intentional fallacy*, because what the author meant is effectively irrelevant once the book is read: its interpretation by the reader *de facto* defines what it says. I therefore review the book as an example of the form to which it appears to aspire (namely a scientific study), not least because the music, in my opinion, doesn't really bear close scrutiny. More on this below.

I have, perhaps provocatively, begun my review with a definition of pseudoscience. Using Wikipedia's feature list, I have labelled points below where this work seems to me to fall into that category, with numbers in square brackets, referring to the list above. I leave the reader to make the final call.

2 Overview

Structurally speaking, a significant chunk of the book consists of a number of journal papers (Cope, 2000, 2002, 2004), placed end to end, with a narrative about creativity, including details about Cope's personal development, to join them together. The parts of the book which have not, it seems likely, been through academic review, are those which give the 'background' on musical-creative systems and, most importantly, the sections on 'association networks' [1b], where creativity is strongly claimed. It's not so unusual to produce a book by splicing papers, but it is appropriate to warn your readers of the fact before they buy, and, much more importantly, to update older work in time for the re-publication. In the current case, many of the claims made (for example, on page 181, that there are 'few published examples of computational learning in music') may have been marginally defensible when the original papers were written, but were certainly not in 2005 (see www.ismir.net for instant refutation of this particular example).

Cope has been working for nearly two decades on a suite of programs which, he claims, 'compose' music (*Preface*) [1a]. Until recently, the aim was to produce music in particular styles (Cope, 1991), most usually classical (in the strict musicological sense of the term, the period roughly from 1740–1820) and Romantic (roughly 1815–1910), and much of the early emphasis was to use the software to support Cope's own compositional activity: his web-site features excerpts from three Grand Operas (his term) written in this way—about composers, and in the style of the composer they are about. Later, the human intervention was reduced, maybe to nothing (Cope, 2000). However, in all of these books, and in the peer-reviewed papers (predominantly in the Computer Music Journal) of which they are mostly composed, I have been unable to find published details (to the extent of reproducibility) of how they work—rather, there are imprecise discussions of representations and rules, filled out with examples that sometimes give us an illusion of understanding what the mechanism does [3].

The question of reproducibility—a basic criterion for publishability in a strictly scientific context—is always a problem for artists who wish to publish the computer-scientific aspects of their work. Harold Cohen, for example, takes the explicit, strong position that the detail of his program, AARON (e.g. Cohen 1979, 1988), is part of his art, and therefore not available for detailed scrutiny—though he does not allow that to prevent him going into detail about how some things can be done, to the level from which they can be reproduced. Cope, on the other hand, does not address this issue of reproducibility (but also does not achieve it), leaving us again unsure as to exactly what his intent is.

For me, the clearest description of how the program *Experiments in Musical Intelligence* works is actually written by Douglas Hofstadter—in one of his two chapters in Cope's (2000) edition of colleagues' discussions of his own work [2b,5a]. I was pleased, therefore, to find that the new volume promises to 'describe', 'explain' and 'define' the author's work (*Preface*)—and equally disappointed to find that, while it does indeed do these things to a far greater extent than previous work, it does not do so to a level even close to one allowing reproducibility of the ideas. Instead, Cope provides the code, written in Lisp. This is certainly admirable, but it is no substitute for a proper statement of algorithms—necessarily, because of the demands of I/O, the code is not declarative; undocumented global variables abound; and, as a whole, it would take several weeks' work to understand the algorithms from the code [1b,3].

3 Definitions of 'Creativity'

The book starts from a definition of 'creativity' which is rather different from those given in the established creativity studies literature (e.g. Boden, 1990, Wallas, 1926). Cope offers the following definition:

> "initialization of connections between two or more multifaceted things, ideas or phenomena hitherto not otherwise considered actively connected" as a cornerstone of this book. Each of the chapters [. . .] will exemplify one or more aspects of this definition.
> (Cope, 2005, p. 26, part in quotes repeated throughout)

Each exemplifying chapter has its own principle, stated at its head, and all the principles are summarised in Chapter 12 (p. 368); however, they do not match the definition. For example, principle 11 says that 'Creativity depends on the integration of its various characteristics into a unified whole'—but the definition is satisfied with mere 'initialization'; principle 12 says that 'Creativity depends on aesthetic values. . .'; but aesthetic values are not mentioned in the definition. Some might say that this complaint is merely being picky—but what is the point of giving an explicit definition if one then contradicts it?

Most definitions of creativity (e.g. Boden, 1998) include some (usually contextual) notion of *value*—a process is not deemed creative unless the artefact produced is valued, nor, indeed, unless it is novel (Bundy, 1994). On the other hand, Cope's definition seems to be very similar to Koestler's (1976) notion of *bisociation of matrices*. Cope, however, does not refer to Koestler, nor to the equally strongly related work on *conceptual blending* (Turner and Fauconnier, 1995) and (re)invents his own terms [1b]: 'inductive association' and 'recombinance', respectively. The implication is that creativity (as defined before) *is precisely* a combination of these things. Koestler, on the other hand, concurs with other authors in deeming awareness of value as a component of creativity itself.

Of course, Cope's definition of creativity is particularly convenient for his style-imitation programs, because the production of such 'associations' and subsequent blending of two musical descriptions is precisely how they work, at least as explained by Hofstadter (Cope, 2000). So by defining creativity thus, in particular without a component of self-evaluation, he neatly (re)defines his work into creativity studies [2a], whereas, formerly, it was concerned with imitation of musical style and not presented as the study of creative systems *per se*.

Bundy (1994) raises the issue of complexity in creativity, and proposes that there is a difference between 'mere novelty' and 'real creativity' based on the complexity of the artefact produced; Wiggins (2006) picks up this argument and suggests that the complexity may lie in the creative process itself, or in the *conceptual space* (Boden, 1990) that the generated artefact inhabits. Cope (2005, p. 27) introduces the notion of 'comtivity' [1b], where something falsely passes for the result of a creative process by virtue of its complexity; but since he does not refer to Bundy's proposal or others related to it, it is hard to see where this idea fits.

The notion of evaluation is problematic here in another way. In the past, Cope has played what he calls 'The Game', presumably, but not explicitly, inspired by Turing's (1950) 'Imitation Game', commonly interpreted (but not explicitly presented by Turing) as a test for intelligence. In this game, audiences are invited to identify computer-composed pieces from human-composed ones—but, we have to note, not arbitrary machine-composed ones, but those chosen by Cope himself [2c]:

> Experiments in Musical Intelligence music, at least the music I have allowed to be performed and published, . . .
>
> (Cope, 2005, p. 256)

In the current volume, this selectivity is much reduced, with the result that some of the musical examples, most problematically the outputs of the improvisation program (Chapter 4), are working at well below the state of the art as exemplified by the work of, for example, Pachet (2002). So as Cope laudibly opens up his proposal to scrutiny, we begin to see the warts—and the drawbacks of a definition of creativity which does not require introspection on quality. Apologising, in this connection, for the weak output of his improvisation program, Cope says:

> Unfortunately, discovering beats in performed music of any style is an enormous analytical problem in itself . . .
>
> (Cope, 2005, p. 115)

and admits

> Of the programs I have created over the years to detect beats in music, few have been successful . . .
>
> (Cope, 2005, p. 116)

A better approach might have been to study the well-publicised work of, for one example of many, Dixon (1997), where very good solutions are available to identify musical beat in audio signals—a problem far harder than beat-tracking in the symbolic representation used here. In this latter case, Sibelius Software's *Sibelius 2* music typesetting package, from more than four years ago, could read expressively performed MIDI files and get the beat tracking right the vast majority of the time.

My point is this: if an author intends to take up his readers' time with descriptions of what he has done, he has an obligation not to waste that time by making claims which are inaccurate—and demonstrably so, with just two minutes' use of Google.

4 Music Analysis

Cope claims two contributions to creativity-related music analysis: the analysis of 'allusions' and a musical-tension-related analysis, which Cope calls 'SPEAC'. Both of these deserve attention; the latter is probably the most useful contribution of the book.

Allusions, Cope tells us, abound in music. This is particularly so in his very liberal definition of allusion, which seems to count any common sequence of two or more notes as an allusion between two pieces, even if the composer did not intend it to be so. Nevertheless, it is analytically useful to be able to detect common material between pieces. Cope gives a program which does this, called 'Sorcerer'. Examination of the implementation shows that all this program does is compare the pieces notewise; it's not surprising, therefore, that when run on large pieces, or large databases of pieces, it becomes far too slow: one is forced to restrict the maximum length of sequence. As far as I can tell from the undocumented code, gaps in allusions are not allowed, so that under this definition, variations on a theme would not count as ‘allusions’. Clausen *et al.* (2000) and Meredith *et al.* (2001) published efficient algorithms to do this kind of analysis (in fact, more, because they are capable of dealing with partial matches). Again, Cope would do well to read the literature.

SPEAC is a symbolic approach to mapping out what is often popularly referred to as the ‘ebb and flow’ of music, and this is indeed a crucial ingredient of any autonomous composition system: modulation of this musical quality is fundamental in producing *good* music, and not just notes, in many, if not most, styles. SPEAC analysis, which relates to uncited cognitive work such as that of Narmour (1990) allows a program to plan these effects into the framework of the music. It is a pity that more of the book was not devoted to the detail of this and the resulting musical constructions, which might have made a real contribution to the literature.

5 AI Techniques and Mechanisms

Cope’s lack of background research is nowhere stronger than in the introductory section on Music Programs and Research, where he examines the broad techniques that have been explored in artificial intelligence research on music. He begins, appropriately, with ‘Rules-Based Programming’. Under this heading, he includes Markov Chains, a probabilistic form of rule-based system. Next, still under the same heading, he generalises Markov Chain application to cases where the probabilities are learned, expanding this into ‘statistical approaches’. All this is fine, in principle, except that the next section is called Data-Driven Programming, and says

> The Markov analysis [...] just described [under “Rules-Based Programming”] is also data-driven, though it differs from data-driven processes in that it uses rules ...
> (Cope, 2005, pp. 62–3)

In other words: $p\wedge\neg p$. The section piles error upon error: Cellular Automata are misclassified as a kind of Genetic Algorithm (pp. 60–1); hidden units in neural networks are apparently so called ‘because their values do not reveal much about their contribution to the network...’ (p. 61)—Russell and Norvig (1995, p. 571) explain the real reason in their AI text-book, which is easily available in libraries throughout the world; and so on. If you find a copy of *Computer Models of Musical Creativity* in your library, I suggest you advise your undergraduate students not to read this section (pp. 57–80), lest they believe it and lose credit in their exams as a result. It’s worth emphasising here that these errors are not matters of subjective opinion, but of defined fact.

On a more positive note, the central idea around which Cope’s ultimate ‘creative’ system, ‘association networks’ is based is much more interesting. This idea, he says, he ‘developed in the mid-1990s’ (Cope, 2005, p. 274) but in fact, these are semantic networks, *pacet* Cope’s incorrect claim that the latter do not come annotated with arc weights: ‘semantic network’ is a general term for a network which represents meaning (Winston, 1984, p. 298), and that is certainly what Cope is doing here. The arc weights are used to represent association strengths as is, for example, common in work on Latent Semantic Analysis (LSA, Deerwester *et al.*, 1990). The fact that the LSA networks are represented as matrices for mathematical purposes is immaterial.

Nevertheless, the key point is that here we are focusing in on what it might mean to (in Koestler’s terms) ‘bisociate matrices’, and Cope shows us how his association network can be trained, like a sort of apprentice Eliza, to return words and phrases which are associated by some kind of lexical proximity in typed input. An operator is provided to mark particular words for some inductive reasoning before the association is applied, allowing limited analogical reasoning. Having been shown some small examples of running the training process of the network with language, we then learn that it can store music as well as words. It’s unclear how this works, but the implication is that the ‘music’ is spelled out as letters; I can’t be sure if this interpretation is correct. Because of the proximity of this ‘music’ to words in the language input with which it is interspersed, associations are made; however, it is entirely unclear if and how musically meaningful associations between musical fragments are made, given this simple level of representation. Further operators are introduced, to allow the user (who takes the role of mentoring interlocutor) to control the range of analogy that can be used (as above) and also to reinforce or suppress the inferences made by the system.

Unfortunately, there is a strong impression of smoke and mirrors with all this. Extraordinary diagrams 'illustrate' these processes, for example in Figures 9.1 and 9.2. The two figures contain transcripts of interactions with the program, accompanied by a 'representation of [the] conversation'. This latter consists of a number of circles most of which are connected together by lines (it's not a fully connected graph, contradicting what Cope says in the narrative). There are no labels on the network, and there are no weights. In short, these diagrams are utterly meaningless and can only be there to create an impression of technical content [1a,1b]; they no more illustrate a scientific process than a picture of flowers on lavender-scented notepaper. We do get to learn the detail in the text, which merely confirms the impression that Cope's 'inductive association' is a very pale reflection of (uncited) LSA, performed on an ill-defined representation. So it is really not clear how this very simple program would produce music of the kind claimed, or of the (albeit not particularly high) quality exemplified [2c,4].

6 The Music

This brings me to '*what matters most*', according to Cope. The book contains many examples of work, which Cope openly acknowledges are hand-picked for quality (see above). Now, of course any researcher selects good examples of his or her work for the purposes of presentation, but rigorous researchers also discuss its failings in context of the literature. Cope does not [2a,2b,2c], for example discuss Bach-style chorale harmonisation without even mentioning the very successful work of Ebcioğlu (1988), which set the standard for this particular task.

The biggest example, billed in the book as 'Beethoven's 10th Symphony, 2nd movement' is available for listening at http://arts.ucsc.edu/faculty/cope/cmmc.html (scroll to the foot of the page and listen to the MP3 version). Aside from the fact that, at least when I listened to it, the recording had one of its string parts played on a percussion channel,[2] making the whole effect quite comical, with snare drum, whistle and even vibra-slap, the musical claim here is laughable in itself. It is not clearly explicit which of Cope's programs produced this piece, but given how it is constructed (by taking the framework of one existing piece, pasting parts of others over it, and distorting them and the framework so that they fit together) it would seem to be the work of *Experiments in Musical Intelligence*, which is odd, since Cope disclaims his previous work on style imitation at the start of the book. In any case, the piece is clearly based on the second movement of Beethoven's fifth Symphony, but a repeated harmonic movement imposed by the program gives it a very odd Baroque feel, like a chaconne. Cleverly, it manages to be recognisable as derived from Beethoven's fifth, without sounding like Beethoven [4]. Because I doubted my own opinion, I asked a colleague, who has a PhD in music, to identify the composer, and he could only (reluctantly) suggest Albinoni, in spite of the fact that he recognised the similarity to Beethoven's fifth as soon as he heard it. Subjective judgements aside, the claim that a piece of music based on the structure of Beethoven's fifth would be a good candidate for Beethoven's tenth is absurd, as listening to Beethoven's ninth immediately shows: Beethoven's style developed and changed considerably during his lifetime, so this notion just doesn't hold water. Either Cope is joking (but it's not very funny), or he's being musically naïve (which seems unlikely given his experience as a composer) or he's assuming his readers are musically naïve (which seems insulting). I don't know which.

7 Conclusion

In summary, I do not believe that this volume contributes anything of substance to the study of computational creativity, nor to the literature on music. The author shows what can only be described as ignorance of the literature in his subject, of which I give examples above. Criticism on these grounds is carefully pre-deflected with a modest apology in the *Preface* to those colleagues whose work he has not been able to include 'because of space limitations'. Failure to cite is one thing, but it is, frankly, shocking that MIT Press should publish an academic volume with such flagrant and basic factual errors in it as I listed above. My suggestion that students be warned is serious: it's likely that less well-informed readers will believe what they see in print in preference to the words of their professors—and thus is inaccuracy, imprecision and ignorance promulgated. What is more, this kind of nonsense brings the whole field of computational creativity into disrepute. It's bad enough that many funding bodies and academic institutions view the study of music as something to be done for fun and not taken seriously,[3] without authors giving the impression that it's done sloppily as well.

Cope is right to propose that machine learning methods do seem to be the future of computational creativity: otherwise, with hand-coded rules of whatever kind, we can never get away from the claim that the creativity is coming from the programmer and not the program. He's also right that musical form is very important in modelling musical creativity; and he has identified important aspects of a creative system for composing music (most notably SPEAC analysis, mentioned before, although this is only a descriptive theory, and not an explanatory one, in the sense of Wiggins, 2007). However, neither of these points is new, nor is either of them surprising to one versed in the relevant literature (note that nearly all the work cited in this review article predates Cope's book by at least four years; and there's plenty more where that came from). Reading this book, therefore, merely deflects the serious student from the extant philosophically and scientifically sound research that it fails to cite.

To return briefly to the question of evaluation: Chapter 12, entitled 'Aesthetics', is not actually about aesthetics, an area which would be overripe for explicit consideration at this point. It is mostly an extended polemic about authors who have criticised Cope's work, under the sub-heading 'Prejudgement'. As the heading suggests, the theme is: 'people don't like this because their minds are closed' [5b], a claim which is not convincingly argued. Later in the chapter, in more positive mood, Cope winds up the book's contribution with three axioms of machine creativity:

(1) Machine programs can create.
(2) The quality of music has nothing to do with who or what created it.
(3) The only limit to what machines can do is the limit of what we as humans can do with machines.

(Cope, 2005, pp. 370–1)

I leave it to the reader to imagine what kind of theory might be based on 'axioms' such as these. Terminology aside, 'axiom' 1 seems to have been demonstrated already, by programs such as HR (Colton *et al.*, 2000), but (even for me, a believer in machine creativity) the case is not made here; 'axiom' 2 is patently false, since the quality of music written by, for example, this reviewer will never be as high as that of music written by, say, Beethoven; and 'axiom' 3 is false too, because machine learning systems in principle allow machines to do more than humans can do themselves, autonomously.

One can do little but regretfully conclude that the reason Cope finds others criticising his work is not because their minds are closed to it, as he says in Chapter 12, but because it is not very good.

Acknowledgement

I am very grateful to Dr Alan Marsden for helpful comments on an earlier draft of this review article.

References

Boden, M. (1990). *The Creative Mind*. London, UK: Abacus.

Boden, M. (1998). Creativity and artificial intelligence. *Journal of Artificial Intelligence*, **103**(2): 347–56.

Bundy, A. (1994). What is the difference between real creativity and mere novelty? *Behavioural and Brain Sciences*, **17**(3): 533–4.

Clausen, M., Engelbrecht, R., Meyer, D., and Schmitz, J. (2000). Proms: a web-based tool for searching in polyphonic music. In *Proceedings of the International Symposium on Music Information Retrieval (ISMIR'2000),* Plymouth: MA.

Cohen, H. (1979). What is an image? In *Proceedings of the 1979 International Joint Conference on Artificial Intelligence*, Tokyo, Japan.

Cohen, H. (1988). How to draw three people in a botanical garden. In *Proceedings of the 1988 Conference of the American Association for Artificial Intelligence (AAAI-88)*, Cambridge, MA: MIT Press.

Colton, S., Bundy, A., and Walsh, T. Automatic invention of integer sequences. In *Proceedings of AAAI 2000*, pp. 558–63. Cambridge, MA: AAAI Press/MIT Press.

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

Cope, D. (2000). *Virtual Music: Computer Synthesis of Musical Style*. Cambridge, MA: MIT Press.

Cope, D. (2002). Computer analysis and composition using atonal voice-leading techniques. *Perspectives of New Music*, **40**(1): 121–46.

Cope, D. (2004). A musical learning algorithm. *Computer Music Journal*, **28**(3): 12–27.

Cope, D. (2005). *Computer Models of Musical Creativity*. Cambridge, MA: MIT Press.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, **41**(6): 391–407.

Dixon, S. (1997). Beat induction and rhythm recognition. In *Australian Joint Conference on Artificial Intelligence* Perth, Australia, pp. 311–20 http://www.elec.qmul.ac.uk/people/simond/pub/1997/ai.ps.gz.

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

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

Meredith, D., Wiggins, G., and Lemström, K. (2001). Pattern discovery and matching in polyphonic music and other multidimensional datasets. In. *Proceedings of 2001 Conference on Systemics, Cybernetics and Informatics*, volume X: Orlando, Florida.

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

Pachet, F. (2002). Playing with virtual musicians: the continuator in practice. *IEEE Multimedia*, **9**(3): 77–82.

Plotkin, H. (1998). *Evolution in Mind*. Cambridge, MA: Harvard University Press.

Russell, S. and Norvig, P. (1995). *Artificial Intelligence – a Modern Approach,* Prentice New Jersey: Hall.

Turing, A. (1950). Computing machinery and intelligence. *Mind*, **LIX**(236): 433–60.

Turner, M. and Fauconnier, G. (1995). Conceptual integration and formal expression. *Metaphor and Symbolic Activity*, **10**(3): 183–203.

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

Wiggins, G. A. (1998). Music, syntax, and the meaning of "meaning". In *Proceedings of the First Symposium on Music and Computers*. Corfu, Greece.

Wiggins, G. A. (2001). Towards a more precise characterisation of creativity in AI. In Bento, C. and Cardoso, A. (eds), *Proceedings of the ICCBR2001 Workshop on Creative Systems*, Naval Research Laboratory, Code 5510, Navy Center for Applied Research in Artificial Intelligence: Washington, DC.

Wiggins, G. A. (2006). Searching for computational creativity. *New Generation Computing*, **24**(3): 209–22.

Wiggins, G. A. (2007). Models of musical similarity. *Musicae Scientiae, Discussion Forum 4a*.

Winston, P. (1984). *Artificial Intelligence* (2nd edn) Boston, MA: Addison-Wesley.

Notes

1 'A proud and meaningless association of human beings' (http://granfalloon.org).

2 Indicative, presumably, that Cope did not check his audio examples any more than he checked his AI background.

3 They're factually incorrect of course. Music is a very serious economic driver: in the UK, for example, income from the music industry well exceeds that from the automotive industry. And the fact that no known human culture exists without music suggests that it is fundamentally important to humanity.