

On the Meaning of Life (in Artificial Life Approaches to Music)

Oliver Bown

Centre for Cognition, Computation and Culture,
Goldsmiths College, University of London,
New Cross, SE14 6NW, UK
o.bown@gold.ac.uk

Geraint A. Wiggins

Centre for Cognition, Computation and Culture,
Goldsmiths College, University of London,
New Cross, SE14 6NW, UK
g.wiggins@gold.ac.uk

Abstract

Artificial life (alife) is of interest to computer musicians due to its generative potential and the potential for producing lifelike behaviours for musical interaction. In this paper we consider how future developments in alife music could have equal bearing on the major themes in alife as on the music it produced. We focus on a discussion of the socio-cultural dimensions of making music with technology and argue that modern popular music making practice outside of individualist academic research projects is an important context for the development of alife music systems. This discussion introduces a number of themes about how computational creativity and human creativity may interact as the field progresses.

1 Introduction

Artificial life (alife) and artificial intelligence (AI) exist as independent subjects: put crudely, life does not require intelligence (the intelligence explored by *good old fashioned AI*) (Brooks, 1990), and intelligence (of that same kind) does not require life. Alife as a whole is unambiguously dedicated to the theoretical study of life, and the experimental study of lifelike systems *in silico*. Any notion of alife music (by which we mean composition and performance using alife systems, rather than the scientific study of music as a system of interaction using an alife methodology) lacks this purity of focus; it is a peculiar hybrid. And yet it is also an emerging field, alongside the use of alife in other arts, which sees great potential in the application of broad computational questions of life within artistic practice, including with respect to the mimicking of human creativity. The purpose of this paper is to untangle the divergent goals of alife and regular music practice and to attempt to focus on a potential common

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

©2007 Goldsmiths, University of London

interest in terms of a *strong* notion of alife music¹.

It is easy to accept a notion of music *inspired by* alife systems, or music that is generated by alife systems, and it is clear that the patterns generated by such systems are likely to have some degree of musical intrigue, both in their local temporal structure and in their ability to generate variation (Miranda, 2001; Berry and Dahlstedt, 2003). But interesting musical structures can also, according to other music practitioners, be made by monitoring atmospheric conditions, cosmic rays, or by recording the sound of wind-induced vibrations on London's Millennium Bridge. Endlessly generative (thus arguably *creative*) systems can also be produced by less exotic approaches than alife, including the simple combinatoric approaches developed by artists like Brian Eno, for whom a specific kind of musical style, ambient music, was required in order to facilitate *acceptable* generative pieces (Eno, 1996) (also see Jem Finer's *Longplayer*, <http://www.longplayer.org>). Alife-inspired or alife-generated music has a place in the context of these existing approaches to music, where interesting natural dynamics and powerful generative methods are appropriated to musical ends. In this paper we will be interested in discussing a coming together of alife and music that is stronger than this, in an attempt to approach long term issues in the development of creative systems.

Alife approaches to music can also fall closely in step with more general AI approaches to music in which the designer's goal is to build a system that achieves some more or less precisely specified musical capacity. Alife is concerned with evolutionary and adaptive systems, and adaptivity is closely associated with the essential nature of life because it defines a system that is able to respond to its environment in a beneficial way, and is therefore able to *survive*. We can use artificial evolution or learning techniques to generate systems that function as musical agents, and this is very interesting because in doing so we have defined an artificial environment, and within that environment the criteria for its inhabitants' survival. But this environment is only partly artificial. The 'artificial' of pure alife is an artificial built exclusively to interrogate the real, but in music it interacts and overlaps with the real; the environment of real musicians making real music. This is

¹By this we mean a strong notion of alife music rather than the notion of strong alife applied to music, although as this discussion unfolds some may have the sense that these are the same.

the starting point for the discussion in this paper. What is *the meaning of life* in these instances?

2 Alife Music in Practice

The first author's interest in alife music stems from an interest in realtime interactive musical agents, in particular agents that are sufficiently complex that they cannot be controlled directly but must be interacted with, where this interaction provokes a sense of engagement that is musically pleasing, that has some of the characteristics of interaction with another live musician. Many commentators have discussed these issues in the emerging domain of live algorithms (*e.g.*, Blackwell and Young, 2004; Collins, 2006).

Strong alife music systems are music systems that music producers and consumers will genuinely feel are autonomous and lifelike, and the above two qualities are, we assume, necessary attributes of such systems. They are also necessarily subjective. A critical question is whether these systems need to emulate human behaviour, or whether there are other modes of behaviour that can evoke a sense of autonomy in musical contexts. Can we have strong alife music whilst bypassing many of the challenges of understanding human intelligence? It is challenging to imagine musical agents that are not human and also do not explicitly mimic human behaviour but that are convincingly autonomous and lifelike: no such thing exists today.

The first author's personal attempts at alife music involve the use of Continuous-Time Recurrent Neural Networks (CTRNNs) (Beer, 1996; Slocum et al., 2000) in live musical performance. The CTRNNs act as simple decision engines in a complex performance patch built in Max/MSP (www.cycling74.com), which pre-processes audio input for feeding into the inputs of the CTRNN and maps the CTRNN's output to generate audio in various ways. So far all of the CTRNNs used are generated by artificial evolution using simple hand-written fitness functions which express the artists own impressions of what would make interesting musical behaviour. This artificial evolutionary environment is very different from the real environment in which the artist actually exists as a musician. Whether explicitly demoing this system or just using it in the context of performance, during performance the system is inevitably squeezed to fit into a musical goal.

Thus it is common to override or tightly constrain the behaviour of a system during a performance. Whilst the simulated environment represents a first approximation of what a musician expects to be good musical behaviour, there is hardly any overlap between the CTRNN's simulated evolutionary environment and the real environment in which the musician acts. The latter is also very many orders of magnitude more complex than the former. It would seem natural to look back over the design of the entire system and ask how one could adjust this design to make the system more successful. This is the point at which a strong alife approach to music must differ fundamentally from other AI approaches.

Consider the iterative process that starts when any music system is first tested in a real performance context. It

would be unusual to have pre-specified quantitative elements to measure the system. Instead it is normal to remain open to the possibility that the system had unexpected positive qualities, even if these are extremely modest qualities. In general we can't predict how people will make use of the system, or what they will make of it, especially if the system is intended to be complex and full of surprises.

Two very simple observations have emerged from working with CTRNNs in a performance context. Firstly, a modest unexpected quality: the activity of the network is valuable in live solo laptop performance even if only to produce loosely synchronised activity, meaning activity which has no precise timing or decision-making demands. This is like having an extra set of hands to control some parameters in a laptop performance. This is a relatively unambitious use of the network, not dissimilar to drawing data from environmental conditions or any of the other examples discussed in the introduction. Our point for the time being is that it is a slight *re-appropriation* in terms of how one conceives of the network and how it is used. Secondly, expectations by third party musical performers about what the network can or should do can become problematic. By disguising the network's behaviour as some relatively direct consequence of a human's activity (the hidden activity of the laptop performer), these expectations are no longer relevant, and this has proven useful. This is not to say that the CTRNN has descended to this level of use, in some cases it has been possible to bring it to the fore and to step away from controlling it, whilst in other cases it is more appropriate to subsume its behaviour under other musical goals and activities. Rather, it proves that the CTRNN serves some purpose as a *mere tool*, with just an inkling of something more worthy of the term alife. This demonstrates what we would call the *cybernetic flexibility* of music as a domain of human activity. In the following section we aim to provide some background to the proposal that this flexibility can and should be captured and used to greater effect.

3 Life

In nature we often view organisms as having adapted to their environments through evolution by natural selection. This is widely understood to be a simplification, albeit one that is highly convenient and often sufficiently accurate. In reality all organisms alter the environment for all other organisms, all evolving, all at once in one large open dynamical system. Lovelock's Daisyworld model provides a definitive proof of concept for this point of view (Lovelock, 1979). The model consists of a planet heated by a sun, and two types of daisies with different heat absorption properties and temperature preferences. Running the model shows the relative populations of daisies stabilising in an arrangement in which the local temperature is optimal for each of the daisy types. Daisyworld's virtual daisies could easily have given the impression to naive observers of having adapted to their environment, but in fact they have altered it through a simple thermo-regulatory process.

The implications of considering the effects of this

true *coevolutionary* process are now beginning to be explored and will continue to drive theoretical biology into the future. For example, Owings and Morton's (1998) new approach to animal vocal communication depicts a natural history full of organisms (assessors) whose naturally evolved senses have become a focus of exploitation by others (managers); owls, for example, that trick badgers into thinking that they are snakes by mimicking snake sounds (the owls don't need to know anything about snakes, or why they make these sounds). This supersedes the older view that implicitly accepted the evolutionary centrality of an animal's senses. Rather, the senses are reconsidered as a context for evolution elsewhere.

Our planet's biodiversity, the celebrated evidence for nature's own creativity, derives from a divergent evolutionary process in which different species react to each other through evolutionary change, rather than from a process, as in most artificial evolution for engineering's sake, which generally aims for convergence on a best solution; nature holds no such requirement on any of its constituent organisms. It is this coevolving process that in some way drives the increase in complexity evidenced in nature's history, where we see not only ingenious solutions to problems, but the creation of new niches and survival challenges themselves, and gentle shifts of focus from one evolutionary process to the next.

How can we begin to bring these kinds of issues into a consideration of alife in the context of music? If an important fact about nature is that it places no strict demands on what life looks like and how it behaves, then how can we reconcile our desire to embed alife in a context in which our own aesthetic requirements are positively stifling?

There appear to be three potential views on this problem. Firstly, the above concerns are excessive and place too high a demand on what we call alife. Engineering-oriented artificial evolution can produce surprising creative solutions to problems and systems that are complex enough that their operation is rendered opaque. Alife is meaningful even in massively constrained, human-centric situations. Secondly, going the opposite way, the strong ambitions of alife music are genuinely flawed; the aliveness and, by implication, the autonomy of a system is hampered by the constraints of an evolutionary context allowing only one course of action: to produce pleasing music or musical behaviour. Natural evolution would not work under such constraints, and so the notion of alife in the context of music is a weak one at best.

The third view, which we prefer, breaks the stalemate of this opposition. It proposes that human musical practice is capable of providing a suitably rich environment for divergent, open-ended evolution to take place. In aspiring to the environmental freedom of nature we must find ways to allow as much variety and flexibility as possible into our demands for artificial musical systems, neither limited to the taste of one human participant, nor static and unresponsive to the actions of agents. In aspiring to the complexity of nature, we must facilitate countless continued repeated interactions between agents and their environments.

These new requirements can be summed up in the proposal that the application of alife in music has suffered

from an individualistic approach. This individualism is manifest in two ways: the individualism of the user, such as in the case of interactive genetic algorithms where a single user is expected to steer an evolving system towards the fulfilment of their musical requirements; and the individualism of the system's purpose, where we assume a musical purpose for our alife system in order to design it.

It may seem that eradicating the second of these two individualisms would ground ones efforts immediately. Is there any sense in producing an alife music system and *then* deciding how to use it? In the following section we suggest that this *is* possible if we look more closely at questions in common music practice and the socio-technological conditions in which music exists.

4 Music Systems, Practice, Sociality and Technology

Strong alife music cannot be about training a system to achieve a pre-specified goal in a pre-specified style. In this case, the less a system's behaviour has been determined by a single individual's expectations, or with respect to a single musical function, the freer it is to take on the properties of an alife music system, rather than an AI or machine learning system. This assertion is strong but not defeatist. However it does jar in various ways with certain instances of the relationship between music, technology, individuals and groups.

To take an example of early work in intelligent music systems, in discussing the role of computers in music performance, Robert Rowe states that "[the elimination of human performers] is undesirable, beyond the purely social considerations, because human players understand what music is and how it works and can communicate that understanding to an audience, whereas computer performers as yet do not." (Rowe, 1993). This is an agreeable statement, except that there is no way of corroborating the simple assertion that "human players understand what music is and how it works". How can one probe this statement further? How do we *know* that it is true, or even what it means? Rowe's "as yet" proposes that computer performers will arrive at a level comparable to human musical understanding. However, the performers that they aim to imitate and perhaps to replace are entangled in a web of relations and social concepts and structures that are suitably versatile as to bring into focus this problem of their understanding, and how it fits with problems of music and understanding in general.

In our opinion the point of view captured here (which we do not mean to associate explicitly with Rowe, but see as a general view about how computational creativity is likely to unfold) does not truly acknowledge the diversity and strength of difference in approaches to music, especially in new social and technological contexts. As Rowe was writing this, the technology behind the tape music approach that he was contemplating was escaping the academic computer music world and initiating the biggest revolution in popular music since rock and roll. Contemporary Western dance music (encompassing genres such as techno, house, garage, drum and bass, hardcore, dubstep and many others) presents problems for any

performance-centric view of music. Much dance music is 'hand programmed' by its producer and at no point during its production or consumption does a performance take place. And yet this is extremely expressive music². If there are clear examples of the widespread social acceptance of non-performed music, even if these examples are based on new technology not possessing a long-standing tradition, then the issue of performance must be understood as a non-essential musical element. For some this may be a relatively minor and unproblematic statement, but it remains an outside point of view and one that is rarely stressed.

But going deeper into the difference between live performed music and studio-based composition, the problem of autonomy in computer music systems becomes centralised through a notion of *editorship*. Imagine a human-edited recording of a piece of computer generated music; let us assume that you enjoy the piece very much, but you have no idea what work was done by the human during the editing phase. If your aim is to judge the musicality of the artificial system, you will find this opacity of presentation naturally quite unsatisfactory: judging a system implies judging it *in action*. For this reason, live performance provides a context in which the evaluation of computer music systems seems to make greater sense. In reality the same problems of editorship still apply. We have described this in the first author's performance work above; the editorship of any software activity generally takes priority. But even in systems that are not tampered with during or after playback, human premeditation and planning are still largely responsible for the ultimate aesthetic and content of the music. Furthermore, through focusing on the musical performance a false boundary arises around the beginning and end of a single performance event. We judge individual human musicians over their careers, and only some musicians in some contexts are valued on the consistent brilliance of their live performances reciting pre-composed music³. More recently, questions in computer music have found themselves inexorably tied up with the booming interest in improvised music, possibly for obvious reasons of suitability – what better test of musicianship than the coming together of the live and the compositional? – but possibly also due to trends in music that are politically broader and more profound, such as discussed by Lewis (2002).

From the point of view of analysing computer musicianship, therefore, it seems more appropriate to state that there is no difference between what is live and what is composed; both can be regarded as performances and neither can be judged for their technological merit from a single instance⁴. And whilst the *humanness* of human

²We say this with some caution because a lot of performance may be associated with the music in clubs or on music videos, and the music may sample other musical performances, and thus disguise human performative involvement in its otherwise mechanical production. Despite this we hold that there are examples of purely non-performative dance music production, as well as electroacoustic and computer tape music.

³Sometimes because they achieve an almost superhuman consistency in their performances, which would make for an ironic criterion for the evaluation of computer performers!

⁴One might argue that an exception lies in cases where, for

musical performance is clearly highly regarded by most people, it is important to acknowledge the possibility that this is not because human musical performance is indelibly written into human music perception, but because we are, for obvious reasons, most familiar with human musical performance, and, as Rowe says, there is presently nothing that approximates it. As is widely observed, individual musical tastes vary to the point of mutual exclusion, and musical styles follow a temporal dynamic that is so rich that anyone should doubt the rigidity of musical tolerance to stop at computer composed music. Meanwhile, socio-cultural factors reify the importance of *human* performance in music: even in studio produced music, such as dance music played by a DJ, the visible act of performance is relished, and the relationship between the audience and that performer is viewed as critical. This highlights the strong relationship between spectacle and musical production, but also a key distinction. Perhaps popular music will always need human performative elements, but this apparently does not place a particularly great constraint on how the music is actually produced.

The discussion of modes of musical production and musical style in this section is aimed at drawing attention to the generally isolated and individualist use of most computer music systems, including alife music systems. In the previous section we argued that individualistic approaches to alife music do not sit well with the understanding that the autonomy of real living systems is contingent on the flexibility with which nature provides what Gibson dubbed *affordances*. This leads to the implication that new approaches to musical production and style are as crucial to the development of concepts of alife in music as are direct advances in the field of alife and more literal developments in applying alife to music in individual situations.

5 How Can We Do Strong Alife Music?

The above discussion would be heading for a completely negative conclusion if it wasn't for the fact that our technological environment is changing the way that people make music, as well as aspects of our social organisation. Most importantly for a notion of alife music, as music producers and consumers increase the degree to which they create music in networked environments, they increasingly contribute to an environment which is genuinely rich in its capacity to generate affordances valuable to evolving software systems. Musicians making music on computers connected to the internet allow for music software that shares information about these various musical contexts. By linking up musical contexts in this way it is possible to see beyond the individualist limitations discussed above, through the creation of a rich and diverse environment. Then no single individual need determine the fate of an alife system, and as a direct consequence of

example, a computer system evaluates a whole piece in order to propose a modification to that piece. This is a process that cannot be placed in a live context because it would require knowledge of the future, suggesting a fundamental difference between live and compositional contexts. But whilst this difference does indeed exist, such cases do not undermine the assertion that computer composition is essentially performative.

this, alife systems need not be subject to any one single functional expectation or interpretation.

Such an observation is not an original contribution by this author. Amongst the various commentators who have discussed the possibility of networked communities of users interacting with communities of software agents, the most significant effort has been made by the Hybrid Society (HS) project (Romero et al., 2003). The HS project aims to explore approaches to artificial evolution in a rich world of interaction generated from a group of individuals interacting over a network, including the internet. They point to the problem of *fatigue* associated with a single user IGA approach, proposing that a multi-user approach is a potential solution to this problem. The HS environment does not make any specific demands about what its agents do, and how they are understood by its human users, except to define a general purpose interaction paradigm. Thus the HS environment provides an appropriate framework for a strong alife music to develop, or at least for interesting provisional research in this domain to take place. However, it does explicitly require that software agents and human users be viewed as equivalent and equal actors in the network, both have the ultimate goal of producing aesthetic artworks.

Although non academic enthusiasts are invited to participate in experiments over the internet, the HS project is strictly executed in an academic experimental manner; participants get involved out of academic interest. This is hardly a surprising state of affairs for current alife music practice, but it is one that tightly limits the potential user-base of any such system, and maintains the separation between real artists working in the real world and the environment in which they interact when they turn their attention to the HS project. For strong alife music the network of potentially interested participants needs to not be restricted by this constraint, and to diversify to the extent of the diversity of current music practice. Also, in order to further consolidate the valuable differences between alife and AI, it is important not to conflate alife agents with humans. A master-pet relationship is a more fitting analogy than one of equivalence.

To expand the user base of alife music software two things need to happen. Firstly, regular music software needs to *go alife*. That is, in normal musical contexts certain elements should be recontextualised as adaptive agents and should be able to gather data that informs the design of new systems, that possibly replace old systems. The most important first step to this is that it continues to behave like regular software. Secondly, music makers need to open up to software that has erratic, unpredictable, idiosyncratic behaviour. The crux of this paper is that a common practice alife music that would satisfy this second condition is feasible, already heavily active, but contingent on the existence of music alife software for its development. It is not necessary to try to define how this common practice alife music would work, but we can consider some questions about it. For this purpose we juxtapose two sci-fi vignettes that capture the essential differences between an imagined alife music and more traditional views of computer intelligence in music:

Sci-fi scenario 1: John is at the concert hall

setting up for his rehearsal. He opens two violin cases and a box containing the Z7 concert-grade violin recital robot. He sets up the Z7's shoulder and arm mechanism on the stage and mounts one of the violins on it. He plugs in the Z7's hardware controller to the shoulder and arm mechanism, and also plugs in a microphone which he points towards himself. He powers up the hardware. A light turns on, red at first, then green after a couple of seconds. The Z7 sounds a pre-recorded note, and they begin tuning their violins. . .

Sci-fi scenario 2: Mark has just got home from school. He logs into his PC, connects his electric guitar to the sound card and starts up AudioLife 5.2. The program asks him whether he would like to load an existing environment from his local machine, or search online for active environments. He chooses to go online, and the software provides a list of current active environments. He browses by category, finally settling for CragFunk, and picks an environment at random from the list. The program asks him if he would like to choose any MIDI or audio files as source material for the software agents, the alternative being that they generate their own material from scratch. . .

The *meaning of life* of the alife music system in the second scenario does not come from the fact that it is *performing live* in a context that is accepted as a site of *real* music. A formal performance (strictly scored or completely improvised) is the tip of the iceberg of a rich inhabitable musical environment. Such a context can be understood to establish a smokescreen between the audience and the performance, and generally imposes strict limitations on the performer's activity (even in the improvised context). Rather it comes from the exploratory day-to-day interactions between human and musical system, which has to be interactive and exploratory for there to be any meaning to the system's action. Thus alife music systems might manifest themselves as small components of existing software systems such as VST plug-ins, plug-ins to music playback software such as Apple's iTunes, and as objects in computer music environments such as Max/MSP (www.cycling74.com), PD (www.puredata.info) and SuperCollider (www.audiosynth.com).

Consider the domain of this exploratory day-to-day musical interaction. How does it differ from context to context? In the vignette the context is that of a child outside of his normal educational routine engaging with some kind of contemporary urban music. He is not a professional. He does not necessarily know what he wants to get out of this interaction, like most children he has a limited sense of what is possible musically, and having grown up with this kind of alife software commonplace he unthinkingly accepts its legitimacy.

In this context there is an important opportunity for the alife system to vary. Each time a download is made from the list it may be the mutated or crossbred offspring of

earlier successful agents. This is different from the variation of a system that is designed to be creative; the user is not the designer of the system, and he does not require that the system is creative and therefore variable, even if he *would* ultimately like the system to pass through long-term changes. The system need not change at all once he has downloaded it. What is important is that when the user does want something to be different they go about finding it in an alife way. This may mean asking for a new variation, as in an IGA, but it could also involve manually tweaking the system as long as the information from this interaction could be interpreted as a significant interaction with its environment, and used to inform later evolution.

There are numerous significant implementation questions surrounding the kind of system that would fulfil this goal. Our concern is only with the context in which this could happen. Could an online multiuser evolving system fit easily with the goals and desires of the people using the system? It would be problematic if users stuck with behaviours they liked and never look for new ones, or if they became frustrated with constantly searching for behaviours which acted in ways they did not understand rather than designing behaviours from known methods. They would soon go back to tried and tested music making, and the system would freeze to a halt. Likewise, as in any evolutionary computing approach, the stagnation of the system itself in local optima is a constant threat to the development of genuinely interesting behaviours, such as the qualities of good alife music systems discussed at the beginning of section 2. Overcoming these obstacles would be an important development in strong alife music, and the design of such systems would ideally be gentle on their demands from end users, or somehow seductive. However, it is interesting to consider the many forms of artistic practice (possibly only in recent history) that depend more on our editorship of existing systems than on a thoroughly creative act. DJing and remixing activities, extending to musical genres such as bootleg and mashup, epitomise this approach to creativity. This is clearly a new cultural paradigm, but it may also be a more explicit expression of what creativity has always been about (*c.f.*, Koestler, 1967; Boden, 1990; Csikszentmihalyi, 1990, 1999). Indeed, various views on the creative process focus on the process of *generate and test*. Thus an optimistic view of strong alife music is that it is actually perfectly suited to our collective creative activities, and blurs the perceived boundaries between the individual as creative system, the society as creative system and software as creative system.

6 Conclusion

Many of the professional music producers of today more often than not learnt their skills outside of the classroom and in the *bedroom studio*, a whole music practice emergent on the technology that was designed around other existing music practices of that time. In this paper we have asked how the strange marriage of artificial life and music *could* come to take on a meaning and significance that truly bears on the principles of alife. We have alluded to the emergence of a new social context that is entire fantasy but with the simple goal of thinking through the

possible ways in which a technological aim and a social practice may come together. It would be wrong to assume that by highlighting this context as a possibility, no matter how theoretically correct it may be, it would be simple to evoke it through some kind of social engineering. In the above discussion we rely on the notion that creative individuals find new uses for existing technology, and there is no reason to believe that the uses they find for alife music systems would pay any homage to the principles of alife. All the same, it is exciting to consider the results of designing multiuser evolutionary systems for popular use that are based on the principle of providing a rich variable evolutionary environment and to study ways in which these systems are taken up, in which case it is vital to acknowledge the role of social trends in the success or failure of such systems, as well as their actual design. This suggests an interesting new direction for alife-based music informatics, which would need to incorporate the analysis of collective human social behaviour in its remit. It also suggests new approaches to the study of computational creativity, in which we sever questions of creativity from intelligence – artistic creativity becomes analogous to the creativity of nature – as well as from the individual – individuals act creatively, but this is only one layer of a greater collective creative process.

Acknowledgements

Oliver Bown's research is supported by a bursary from the Department of Computing, Goldsmiths College. We would like to thank Alice Eldridge for valuable discussions and the three anonymous reviewers for their comments.

References

- Beer, R. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 421–429. MIT Press.
- Berry, R. and Dahlstedt, P. (2003). Artificial life: Why should musicians bother? *Contemporary Music Review*, 22(3):57–67.
- Blackwell, T. and Young, M. (2004). Self-organised music. *Organised Sound*, 9(2):137–150.
- Boden, M. (1990). *The Creative Mind*. George Weidenfeld and Nicholson Ltd.
- Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, 6:3–15.
- Collins, N. (2006). *Towards Autonomous Agents for Live Computer Music: Realtime Machine Listening and Interactive Music Systems*. PhD thesis, Centre for Science and Music, Faculty of Music, University of Cambridge.
- Csikszentmihalyi, M. (1990). The domain of creativity. In Runco, M. and Albert, R. S., editors, *Theories of Creativity*. Sage Publications.

- Csikszentmihalyi, M. (1999). Implications of a systems perspective for the study of creativity. In Sternberg, R. J., editor, *The Handbook of Creativity*. CUP.
- Eno, B. (1996). *A Year With Swollen Appendices*. Faber and Faber.
- Koestler, A. (1967). *The Ghost in the Machine*. Hutchinson and Co.
- Lewis, G. E. (2002). Improvised music after 1950: Afrological and eurological perspectives. *Black Music Research Journal*, 22:215–246.
- Lovelock, J. (1979). *Gaia. A New Look at Life on Earth*. OUP.
- Miranda, E. (2001). *Composing Music with Computers*. Focal Press.
- Owings, D. H. and Morton, E. S. (1998). *Animal Vocal Communication: A New Approach*. Cambridge University Press.
- Romero, J., Machado, P., Santos, A., and Cardoso, A. (2003). On the development of critics in evolutionary computation artists. In *Applications of Evolutionary Computing: EvoWorkshops 2003: EvoBIO, EvoCOP, EvoIASP, EvoMUSART, EvoROB, and EvoSTIM, Essex, UK, April 14-16, 2003. Proceedings*, volume 2611/2003 of *Lecture Notes in Computer Science*, pages 559–569.
- Rowe, R. (1993). *Interactive Music Systems*. MIT Press.
- Slocum, A., Downey, D., and Beer, R. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S., editors, *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*, pages 430–439. MIT Press.