

Creative Agency: A Clearer Goal for Artificial Life in the Arts

Oliver Bown and Jon McCormack

No Institute Given

Abstract. One of the goals of artificial life in the arts is to develop systems that exhibit creativity. We argue that creativity *per se* is a confusing goal for artificial life systems because of the complexity of the relationship between the system, its designers and users, and the creative domain. We analyse this confusion in terms of factors affecting individual human motivation in the arts, and the methods used to measure the success of artificial creative systems. We argue that an attempt to understand *creative agency* as a common thread in nature, human culture, human individuals and computational systems is a necessary step towards a better understanding of computational creativity. We define creative agency with respect to existing theories of creativity and consider human creative agency in terms of human evolution. We then propose how creative agency can be used to analyse the creativity of computational systems in artistic domains.

1 Introduction

1

Both artificial intelligence (AI) and artificial life (Alife) have been used to study artistic creativity and to create new forms of art. Traditionally, AI has focused on the artificial simulation of human intellectual capacities, whereas Alife takes its inspiration from the creative power of nature through processes such as self-organisation, natural selection and autonomy. The study of Alife therefore holds special significance for the arts due to its inherent concern with creativity beyond human agency, paying special attention to systems that exhibit the emergence of new, higher-level primitives in a system *de novo* [1]. Despite these differences of focus, in both approaches artificial creativity is a commonly stated goal, whether represented as a means for better understanding human creativity, creativity in general, or towards new systems for artists. But although the intent is clear, a perspicuous definition of this goal or means of objective measurement remains conspicuously hazy. As Saari and Saari put it, “Creativity is fascinating! We know so much about the topic without having the slightest idea what it is” [2, p. 79]. Such expressions of frustration are not uncommon in the literature, and indeed, creativity shares in this a certain property with

¹ Draft, to appear in Proceedings of European Conference on Artificial Life, Budapest, 2009.

life: we are sure that we know it when we see it, but we don't quite know how stable such certainty might remain in the light of advances in the sciences of the artificial.

Our motivation is a lack of focus on *agency* in the literature on creativity. We argue that a better understanding of creative agency will help clarify the goals of achieving creative behaviour in computational systems.

2 Defining Creative Agency

A typical definition of creativity (e.g. [3]) is as follows:

Definition 1. *A system is creative if it produces novel and valuable (appropriate, useful) output.*

Understandably, the novelty and value of the output of a system have been predominant areas of interest in the literature on creativity. In this paper we turn to the process of production itself: the relationship between subject (the system) and object (the output). We address this relationship in terms of what we call *creative agency*: the extent to which the subject is responsible for *producing* the object.

Definition 2. *The creative agency of a system is the degree to which it is responsible for a creative output.*

Identifying creative agency therefore involves the (apparently subjective) evaluation of responsibility. It is not the output itself that we are interested in, but the creativity of the output, in other words, the intangible qualities of novelty and value. Thus a master artist could employ skilled students to create a work, not once touch the work, but still be attributed with the creative agency of the work. By the same reasoning, a wealthy patron commissioning such a work could take some credit for making the work come about, but their choice to employ a reputed artist would be to borrow already existing creativity. Creative agency clearly relates to concepts of autonomy and adaptiveness. However, we prefer not to pin creative agency to any of these terms explicitly, because, as discussed in the following section, it can be associated with quite diverse structures and processes, including the process of natural evolution. A natural process can be thought of as an *invisible hand*, but is not a tangible entity in the sense in which we think of an autonomous system. Adaptiveness implies adaptation to fulfil a need, so again, natural evolution's blind creative 'search' is not itself adaptive (although populations certainly adapt to environments). While it is tempting to think of creative agency as being directly dependent on the extent to which a system is capable of appropriately evaluating the output (following the reasoning that responsibility for a creation requires awareness), this needn't be the case.

In computational creativity, the problem of creative agency is often taken as being of secondary importance to the novelty and value of the output produced

by a system. A lack of attention to the nature of creative agency is common when discussing creativity in humans, because it is generally taken as given that humans are the only kind of creative agent we need consider. In the case of computational creativity, however, this can be a source of opacity, since we cannot directly translate the notion of creative production that applies to humans straight onto computational systems. Computational systems have a completely different relationship to their environments from people. Not least, they are invariably brought into the world by human design. By highlighting this relationship, computational creativity throws into light the problem of creative agency not only in computational systems, but also in human and natural systems.

An understanding of creative agency has a clear role in the evaluation of creativity in computational systems. It would be misleading to discuss the output of a system without considering the relative contribution of the elements involved in producing that output. Systems that exhibit a low degree of creative agency make a smaller genuine contribution to the novelty and value of the output they are involved in producing; in such cases the creative agency should instead be attributable to the designer of the system. A system that has a high degree of creative agency, on the other hand, should have a greater claim to the novelty or value identified in any output produced by that system. If the output is indeed creative (to be determined separately) then by virtue of its greater contribution to that output, the system itself can be deemed creative. In short: novelty and value that cannot be directly attributed to the computational system should have no weight in supporting claims about the creativity of that system.

We can think of the assignment of creative agency to computational systems as akin to assigning royalties to a collection of artists who collaborated on a creative work – the greater their original contribution to the the output, the higher the attribution of creative agency.

A simplified representation of the problem of creative agency is shown in Figure 1. Agency and creativity are placed on distinct axes (without necessarily implying that they are independent), and we consider two hypothetical computational systems. System A has a high degree of agency but does not produce particularly creative output, whereas System B’s output is highly creative even though the system itself is not particularly responsible for the creativity of that output. The diagonal line represents a hypothesised limit of current systems. At present, designers of computationally creative systems are forced to find compromises between systems of type A and systems of type B, but one of the ultimate goals of computational creativity is to find systems that exhibit the agency of System A, but with true creative output as in System B.

Thus, on the one hand we can create systems that act as generative proxies for our own creativity, but that are not themselves the source for the creativity evident in the output (like System B). An example of this is Harold Cohen’s *Aaron* painting system, which embodies a detailed set of rules encoded by Cohen[4]. *Aaron* acts as a sort of cognitive tool for Cohen, mass-producing variations on an artistic theme overnight, for Cohen to sift through in the morning. *Aaron* extends Cohen’s creative output in novel ways, but it remains low on the agency

axis: it is clear that Aaron’s agency is limited, as the program inhabits a tightly restricted space of picture styles and is in no way adaptive to its environment. On the other hand, we can attempt to increase agency by developing systems that do not explicitly embody our own design decisions. This is a challenging problem and, according to the hypothesised limit in Figure 1, is likely to render the output of the system less creative. It is particularly important to note that whilst the former approach can be clearly stated as a problem in cognitive science, the latter spills open into social and artistic domains of enquiry by prompting discussions of authorship and the social reception of creative work. The recent *DrawBots* project has started to address such questions by attempting to produce a robot artist that was the true author of its artistic output, but achieves only a ‘minimal creativity’ [5].

In the remaining sections, we will expand on the relationship between creative agency and social behaviour. This leads us to a view of creativity as a process operating over a variety of distinct levels, between which value provides an important structural relation.




Fig. 1. Graph representing a hypothesised limit to the combination of agency and creativity in current computationally creative systems.

3 Creative Agency in Humans and Natural Systems

The human species is eminently capable of introducing novel structures into the world, and the same is patently true of nature. Adopting a universal approach, the Alife perspective on creativity might well question the individualistic origins of much creativity. Since we would not expect to identify the creative agency of nature in the random mutations of genetic variation, it would be reasonable

to suspect that the creative agency of culture as a whole is greater than the sum of the creative agencies of individual humans. A number of contemporary trends in human evolutionary theory have expounded this argument. Meme theory, for example, proposes that cultural behaviour can be explained in terms equivalent to genetic theory, by positing the meme as an abstract cultural *replicator* [6]. Memeticists are at pains to stress that this is *not* merely a theoretical construction by analogy, but constitutes a proper Darwinian formulation of culture. Thus beyond their replicating qualities and a process of selection, memes and genes needn't have anything in common; meme propagation should be understood as having radically different properties to gene propagation, including the potential for horizontal and oblique transmission between carriers, different rules for survival and extinction, a completely different mechanics of variation and recombination, and completely different rules for coinhabitation in the same carrier. The meme's mechanism for variation and recombination is, in fact, what we'd typically think of as human creativity, so meme theory provides one way of thinking about human creativity as a subroutine in a greater algorithm for creativity at a cultural level, albeit far more powerful than random mutation. As Newton put it, he could see so far only because he was standing on the shoulders of giants.

A related approach to cultural evolution is that of Boyd and Richerson, who hypothesise that the human capacity for culture is a powerful adaptive product of the benefits to the individual of social learning, but with this generally beneficial adaptation comes the potential for much irrelevant behaviour to become amplified and embedded in stable structures [7]. These structures are what we refer to as *culture*, and can be viewed as a kind of evolutionary baggage; much of it may well be adaptive, but not *necessarily* so. Some have suggested that aspects of culture as powerful as language could indeed be understood as *parasitic* [6]. Boyd and Richerson liken cultural evolution to the runaway evolutionary dynamics identified in sexual selection. Henrich and Gil-White extend Boyd and Richerson's model to incorporate a theory of prestige as a mechanism that could drive the emergence and reinforcement of exotic cultural behaviours in a tight positive feedback loop: as individual fitness becomes increasingly dependent on socially determined factors, human behaviour becomes increasingly de-coupled from its ecological foundations, and self-determining [8].

4 The Importance of Levels

Such theories of bio-cultural coevolution are relevant to our discussion of creative agency because they offer a compelling explanation for the emergence of those creative domains that appear to have only a tentative functional relationship to human survival, namely, the arts and music [9]. Although remaining a speculative field, it points to the need for a multi-layered model of creativity that unifies individual creativity with super-individual cultural processes responsible for driving the emergence of creative domains themselves. The need for distinguishing creativity on different levels is also driven at the sub-individual level

by the perspective in cognitive science typified by Andy Clark's *Extended Mind* hypothesis, which postulates that cultural artefacts offer cognitive cybernetic enhancement [10]. According to this point of view human creative agency is already a highly distributed network of elements with human brains at the centre (for the time being), and computational systems can be understood as active contributors to the creative processes of individuals. This distributed model complements the point of view that the brain itself is a distributed set of functional units, as typified by the *Swiss Army-knife* model of mental modularity, proposed by Barkow, Cosmides and Tooby as a central foundation for the empirical study of evolutionary psychology [11]. Theorists of creativity have revisited this point of view in different ways, most notably Arthur Koestler's metaphor of the *Janus-Faced Holon* [12].

Considering creative agency allows us to distinguish between creative processes operating at different layers. As the above discussion reveals, carving the world into distinct creative agents at different levels is an inherently difficult task, in much the same way as identifying units of selection in evolution, or trying to state precisely what a meme is. Accepting creative agency on different levels complicates the centrality of individual human creative agency, but it offers a better distinction on forms of creativity in both cultural processes and computational systems. We turn now to consider how traditional views of creativity work in this multi-layer model.

5 A Layered View of Traditional Dimensions of Creativity

Creativity theorist Csikszentmihalyi sees novelty and value as the cornerstones of creativity: something can't be creative if it has been done before, but mere banal or inappropriate novelty does not qualify as creativity [13]. Creativity occurs when people find new ways to achieve goals. For Csikszentmihalyi, the problem of defining value in the attribution of creativity necessitates a view centred on the embeddedness of individuals within creative domains:

There is no way to know whether a thought is new except with reference to some standards, and there is no way to tell whether it is valuable until it passes social evaluation. Therefore, creativity does not happen inside people's heads, but in the interaction between a person's thoughts and a sociocultural context. It is a systemic rather than an individual phenomenon. [13, p.23]

Csikszentmihalyi therefore defines a creative person as "someone whose thoughts or actions change a domain, or establish a new domain" (p28). Since modifying a domain influences the way that domain will respond to future potential creativity, individual and domain are strongly interdependent. Csikszentmihalyi's characterisation of the relationship between individual and domain extends naturally to a general relationship between creative agency at different levels, that has an unmistakably Darwinian, or perhaps more appropriately

ecosystemic, feel. The creativity of certain individuals is determined by processes occurring at a higher level (the creative domain), mediated by the generation of a system of value. This is Darwinian in that a higher level process selectively filters elements being produced and reproduced at lower levels. The system of value, like the Darwinian concept of fitness, is implicit and mostly revealed in hindsight.

Novelty itself must also be seen as domain specific if it is to have any non-trivial meaning. Trivially, everything that is different is novel. It is less trivial, and far more meaningful, to measure the *degree of novelty* of things. But measurement occurs in a metric space, and metric spaces are not real things, but are constructed by perceiving agents. This is not a problem for creativity *per se*: novelty is our evidence for creativity, but creative systems don't need to recognise novelty to be creative. This suggests that novelty-seeking alone may have little functional utility. Some human cultures, such as Western industrialised society, seem to have fostered neophilia, forging an inherent link between novelty and value. It is far from evident that this has any functional utility, however [14].

Boden discusses the cognitive requirements for humans to find new ways to achieve goals, distinguishing between three kinds of creative process: combinatorial creativity is the combination of existing elements to create new elements; exploratory creativity is search through an existing conceptual space; and transformational creativity is the transformation of an existing conceptual space. A problem for the precise application of Boden's theoretical work has been the formulation of what these conceptual spaces actually are [15], particularly with respect to understanding how transformational creativity differs from exploratory creativity [15, 16]. Viewing creativity at multiple levels allows us to hypothesise that transformational creativity is really a process occurring at the higher cultural level, for example in the way described by Kuhn in his theory of scientific revolutions [17].

6 Categorising Agency in Computational Creativity

The multilevel approach to creativity helps to identify three distinct ways in which computational systems can exhibit creative agency. The first is by actively contributing to the creative agency of individual humans, as an active component in a distributed creative process. Most computationally creative systems to date fit this category, although they may be at odds with their designers' original goals of establishing human-like creativity. This adheres to the extended mind perspective that our individual creativity is already highly distributed and enhanced by cultural artefacts, some of which may be computational systems performing complex tasks. We already use computers creatively, but their role in our individual creativity is creeping towards an increasingly active status. Programs like Cohen's *Aaron*, and the general increase in popularity of generative art demonstrate how this shift is taking place.

Disappointingly, the predominant tool of Alife-based art – the interactive genetic algorithm (IGA) [18] – has had limited success as a tool for enhancing

creativity. The IGA aims at fixing the problem of formally defining complex human aesthetic preferences by letting humans take the place of the fitness function, but this arguably leads to a poor creative partnership where both user and algorithm assume roles of little creative agency. The genetic algorithm is passive in that it relies on the user for the crucial step of selection, but the user is rendered passive by being unable to control the long-term course of evolution or the underlying structure of the developmental process. Nevertheless, interactive genetic algorithms are beginning to emerge in commercially available creative software where their use makes sense. Dahlstedt’s *Mutasynth*, for example, assists a user to search a vast space of possible synthesiser sounds using an IGA with visual representation of synthesis parameter space [19]. Anyone who has played with a synthesiser will be familiar with the mild sensation of blind search already inherent the mapping from parameters to sounds.

The second approach is to consider how computational systems can fit into existing processes at the higher cultural level. These systems need to identify how individual interactions lead to social structures and cause cultural change. Pockets of research have been conducted in this area, which, as discussed, span a variety of disciplines. The *DrawBots* project attempted to accentuate the social construction of a robotic art system’s creative agency as far as possible by allowing its creations to be exhibited in an art gallery without human intervention, illustrating the potentially vast variety of ways creative agents might manipulate creative domains. This includes the potential circularity hinted at by Henrich and Gil-White [8], that perhaps the legitimisation of the art gallery is enough to *make* the work acceptable to a receptive audience. That said, if, in hindsight, the robot did have an impact on its creative domain in this way, the sticky problem is that the agency of this particular act (putting the work into the gallery) falls yet again to the human agent that curated the event.

Romero, Machado and Santos’ ongoing *Hybrid Society* project aims to build a virtual social system coinhabited by human and computer artists, all operating as both producers and critics and interacting in social networks such that the real artistic value systems of the humans influence the world of the artificial agents [20]. In principle, in such an environment (as with DrawBots), agents may potentially influence the creative domain of human participants. By Csikszentmihalyi’s definition, nothing could provide a better indication of creative agency than this.

Earlier Alife style models, based completely *in silico* (e.g. [21–23]) have already established the potential of exploring basic cultural or bio-cultural dynamics using multi-agent systems, yet it is hard to ground those dynamics in a way that produces anything we would recognise as creative (novel, yes, but of any aesthetic interest, no!). This overlaps smoothly with our third suggested approach, which is to work out how to exploit the creative potential already under investigation in *in silico* research in Alife, but in artistic domains. A pioneering example of such research is the Italian composer Agostino di Scipio’s musical performances, which work by building *sonic ecosystems* that transform the latent sound of the performance space into musical works using a series of

complex variations on the process of audio feedback [24]. Di Scipio’s insight is to begin with the medium that he is interested in, and construct complex networks of processes within that domain (sound itself). In other artistic domains, elements from Alife can be used more literally, such as Jon McCormack’s installation, *Eden*, which presents a population of artificial learning agents whose environment is ‘fed’ by the presence of audience members, who are lured to stay in the installation space by the agent’s ability to create interesting music [25]. *Eden* creates an evolving symbiotic relation between the audience and artificial agents. In these domains, creative emergence can occur that is inherent to the environment defined by the work, and as such, the works do achieve creative agency, without conflicting with the creative achievements of the artists involved in making them. This is what Whitelaw refers to as *metacreation* [26].

7 Conclusion

In this paper, we have argued, for purely practical reasons of evaluation, the need to consider the creative agency of systems that are involved in producing a creative output. Although we believe that this focus will help to clarify the goals of computational creativity and the potential role of Alife in this domain, our contribution does not take the form of a mathematical definition of creative agency which could be easily applied by researchers to various creative systems. Instead, it appears necessary that assigning creative agency will continue to be a subjective matter based on disparate evidence. Our goal has been to attempt to form an appropriate perspective with which to simultaneously view creative processes in nature, human culture, individual human behaviour and existing computationally creative systems. We have argued for a perspective that recognises creative agency and the role of value in mediating between levels in a hierarchy of creative processes. This replaces the dominance of the human individual as the exemplary creative agent with a more distributed set of interacting elements into which computational systems can more easily situate themselves. We propose that this clarifies the potential creative role of Alife systems in the cultural domain of the arts. Such a perspective can ultimately lend itself to more detailed numerical analysis of creativity, however, further discussion combining sociological, philosophical and Alife-based reasoning will be needed before this can be achieved.

References

1. Cariani, P.: Emergence and Artificial Life. In: Artificial Life II, SFI Studies in the Sciences of Complexity. Volume 10. Addison-Wesley, Redwood City, CA (1991) 775–797
2. Saari, D.G., Saari, A.L.: Toward a mathematical modelling of creativity. In Andersson, A.E., Sahlin, N.E., eds.: The Complexity of Creativity. Kluwer Academic Publishers, Netherlands (1997) 79–103
3. Boden, M.: The Creative Mind. George Weidenfeld and Nicholson Ltd (1990)

4. McCorduck, P.: AARON's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen. W. H. Freeman and Co. (1990)
5. Bird, J., Stokes, D.: Evolving minimally creative robots. In Colton, S., Pease, A., eds.: Proceedings of The Third Joint Workshop on Computational Creativity (ECAI '06). (2006) 1–5
6. Blackmore, S.J.: The Meme Machine. OUP, New York (1999)
7. Richerson, P.J., Boyd, R.: Not by Genes Alone: How Culture Transformed Human Evolution. University of Chicago Press, Chicago, IL, US (2005)
8. Henrich, J., Gil-White, F.J.: The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behaviour* **22** (2001) 165–196
9. Bown, O.: Theoretical and Computational Models of Cohesion, Competition and Maladaptation in the Evolution of Human Musical Behaviour. PhD thesis, Department of Computing, Goldsmiths College, University of London (July 2008)
10. Clark, A.: Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence. Oxford University Press (2003)
11. Barkow, J.H., Cosmides, L., Tooby, J.: The Adapted Mind: Evolutionary Psychology and the Generation of Culture. OUP, New York (1992)
12. Koestler, A.: The Ghost in the Machine. Hutchinson and Co., London (1967)
13. Csikszentmihalyi, M.: Creativity: Flow and the Psychology of Discovery and Invention. Harper Collins, New York (1996)
14. Gray, J.: Straw dogs: thoughts on humans and other animals. Number 246. Granta Books, London (2002)
15. Wiggins, G.A.: Towards a more precise characterisation of creativity in AI. In Weber, R., von Wangenheim, C.G., eds.: Case-Based Reasoning: Papers from the Workshop Programme at ICCBR'01, Washington, DC: Naval Research Laboratory, Navy Centre for Applied Research in Artificial Intelligence (2001) 113–120
16. Thornton, C.: How thinking inside the box can become thinking outside the box. In Cardoso, A., Wiggins, G.A., eds.: Proceedings of the 4th International Joint Workshop on Computational Creativity, Goldsmiths, University of London (2007) 113–119
17. Kuhn, T.S.: The Structure of Scientific Revolutions. 3 edn. University of Chicago Press (1996)
18. Takagi, H.: Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE* **89** (Sep 2001) 1275–1296
19. Dahlstedt, P.: A mutasynt in parameter space: interactive composition through evolution. *Organised Sound* **6**(2) (2006) 121–124
20. Romero, J., Machado, P., Santos, A.: On the socialization of evolutionary art. In Giacobini, M., Brabazon, A., Cagnoni, S., Caro, G.A.D., Ekárt, A., Esparcia-Alcázar, A., Farooq, M., Fink, A., Machado, P., McCormack, J., O'Neill, M., Neri, F., Preuss, M., Rothlauf, F., Tarantino, E., Yang, S., eds.: *EvoWorkshops*. Volume 5484 of *Lecture Notes in Computer Science*, Springer (2009) 557–566
21. Saunders, R., Gero, J.S.: Artificial creativity: Emergent notions of creativity in artificial societies of curious agents. In: *Proceedings of Second Iteration*. (2001)
22. Miranda, E.R., Kirby, S., Todd, P.M.: On computational models of the evolution of music: From the origins of musical taste to the emergence of grammars. *Contemporary Music Review* **22**(3) (September 2003) 91–111
23. Werner, G., Todd, P.M.: Too many love songs: sexual selection and the evolution of communication. In Husbands, P., Harvey, I., eds.: *Proceedings of the Fourth Eu-*

- ropean Conference on Artificial Life, Cambridge, MA: MIT Press/Bradford Books (1997) 434–443
24. di Scipio, A.: Sound is the interface: from interactive to ecosystemic signal processing. *Organised Sound* **8**(3) (2003) 269–277
 25. McCormack, J.: Eden: An evolutionary sonic ecosystem. *Advances in Artificial Life, Proceedings of the Sixth European Conference, ECAL LNCS 2159* (2001) 133–142
 26. Whitelaw, M.: *Metacreation: art and artificial life*. MIT Press, Cambridge, Mass. (2004)