

Is it Time for Computational Creativity to Grow Up and start being Irresponsible?

Colin G. Johnson
School of Computing
University of Kent
Canterbury, Kent, UK
C.G.Johnson@kent.ac.uk

Abstract

A recent definition of computational creativity has emphasised that computational creativity systems should “take on certain responsibilities” for generating creative behaviour. This paper examines the notion of responsibilities in that definition, and looks at a number of aspects of the creative act and its context that might play a role in that responsibility, with an emphasis on artistic and musical creativity. This problematises the seemingly simple distinction between systems that have responsibilities for creative activity and those which support or provide tools for creativity. The paper concludes with a discussion of an alternative approach to the subject, which argues that the responsibility for creative action is typically diffused through a complex human/computer system, and that a “systems thinking” approach to locating computational creativity might ask better questions than one that tries to pin creative responsibility to a particular agent.

Introduction

A recent paper by Colton and Wiggins (2012, p21) gives a succinct definition of computational creativity as “the philosophy, science and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative. Compared to earlier attempts to define this area, this definition is notable because it does not define computational creativity with regard to human creativity.

By contrast, earlier definitions have been grounded in comparisons with human creative behaviour. For example, Ritchie (2007, p69) grounds his list of criteria for attributing creativity to a computer program thus: “A central assumption here is that any formal definition of creativity must be based on its ordinary usage; that is, it must be natural and it must be based on human behaviour.”. Furthermore, an earlier overview by Colton, de Mántaras and Stock (2009, p11) begins with the statement that, “At its heart, computational creativity is the study of building software that exhibits behavior that would be deemed creative in humans.”.

In this paper I will explore a specific phrase in the definition— “taking on particular responsibilities” which is the main difference with previous definitions. I would like to explore where these “particular responsibilities” might sit

in the creative process, and how the use of computers might change our idea of where that responsibility might sit. In particular, my focus will be on artistic and musical creativity, though there may be implications for other creative areas.

Who/what is “responsible” for a particular creative artistic act? We can argue that there are a number of things that share this responsibility (here we frame these in the context of a human artist):

- The artist themselves, their actions and patterns of behaviour.
- The artist’s motivation to create the work.
- The background knowledge that the artist has acquired through life, which reflects their general cultural background and specific things that they have encountered or learned.
- The context in which they are making the work.
- The materials that they are using to make the work. In particular, the “resistance”, “grit” and “grain” offered by some materials, which can provide new material that can be serendipitously exploited by the artist.

It is commonly seen as the first of these as the action that takes on “responsibility” for the artistic creation. However, when we try to pin down why this is so, we might start by arguing that had the artist not decided to carry out that particular behaviour, to decide not to create that particular work, then the work would not exist. But, the same argument can be applied to the other items on the list: had the artist not had the relevant background knowledge, or had the material worked in a different way, and so on, the work would not have been capable of being created. We can, of course, take this argument to ludicrous extremes—part of the responsibility for the art being the artists own existence, etc.

Indeed, this is not just an intellectual exercise; determining the responsibility (or credit) for a creative act is important for legal arguments concerning intellectual property rights. McGregor (2014) has recently argued that the legal arguments around creativity might provide a framework for considering computational creativity; along similar lines, Koza (2010) has argued for the use of patentability as a criterion to determine when an AI system is creating artefacts that require “human-competitive” levels of intelligence.

We might stop at a “proximate” cause as being the primary point at which the responsibility lies. But, what is the proximate cause? We might argue that the immediate actions of making the art are the artist’s behaviours, in putting pencil to paper in a particular way. But, even at such a proximate level, we can see that that activity interacts closely with the artist’s motivation, and that during the time-span of creating even a single, simple piece of work there might be a complex interaction between motivation and action.

So, where is the computer in all of this? I have argued elsewhere (Johnson, 2012) that computational creativity research has focused too much on the role of the individual creator, favouring the view of the creative “romantic hero” over forms of creativity that are based on collaboration or the mediation of interaction. In this paper I would like to argue further that the nature of computer-grounded artistic creativity makes assigning this responsibility even harder than it would be for traditional artforms.

The remainder of this paper splits into three sections. The first is concerned with the role of materials, and in particular whether computational artistic and musical materials present a particular challenge for making the distinction between passive tools/materials and active agents to which creative responsibility can be ascribed. The second is an examination of context and background, and considers, through examples of search-based art and semantic mass, whether these can be considered to have any responsibility for the creative action. Finally, a concluding section examines whether a better way to examine this is through a “systems thinking” view, rather than a view based on the notion of responsibility.

Materials

What is an artistic material, or an artistic tool, and how does it differ from something that plays a collaborative role in artistic creation? For traditional artworks, the distinction is clear. For example, an artist uses a tool such as a pencil to create their art, a musician uses an instrument to create a piece of music. Part of an artistic training is to learn to “master” such tools; to learn how to realise artistic intents through the coordinated use of perception, thought and the manipulation of tools. But, even at the level of simple physical tools, there is some level of interaction between the tool and artistic creation—part of the study of a particular artistic medium is learning its constraints, and learning how to make adaptations when a particular intentional action does not realise the intended aim.

Certain computational artistic media and tools make this distinction between passive tools and media that are manipulated by an artist or musician, and participants that take an active (creative?) role in the artistic creation. Rowe (1993) has discussed a continuum of interactive computer music systems, ranging from simple action-response systems where a performer makes a physical gesture and generates a consistent sound, to systems which listen and make sound as an autonomous and equal participant in a musical interaction. An example of the latter might be the *Voyager* system (Lewis, 2000); these ideas have been taken further by Paine (2002).

We would probably consider something to the latter end of the continuum to sit comfortably within definitions of computational creativity such as that discussed at the beginning of this paper. Whilst a system of that kind might always perform within an interactive context with other (human or computer) performers, it is “responsible” for holding up its own end in the music being produced, and “creating” music that is sensitive to the current situation. There are multiple “responsible” agents involved, and nothing playing the role of “mere tool”.

But, when we look at systems towards the middle of that continuum, the allocation of creative responsibility becomes murkier. For example, consider the LIES system by Sanfilippo (2012). This consists of a number of acoustic feedback loops, which initially create sounds by creating positive feedback cycles that can start from tiny fluctuations in the performance environment. These are modified by a large number of digital filters and feedback networks. The performer interacts with this system by adjusting the parameters of the various filters and intensity of the feedback system.

What is *responsible* for the final creative output in this system? The interaction between human and machine is complex and at times incomprehensible to the human; the performance mode is one where the human sometimes tries to control the sound being generated to bring it into line with a desired sound (sometimes successfully, sometimes not), sometimes just letting the sound unfold without interference, and sometimes to explore the effect of parameter changes with, depending on context, a greater or lesser understanding of the likely effects. Certainly, the system generates a decent amount of the creative material here, with the human sometimes (importantly—not all of the time) being unable to shape the system’s outputs in any comprehensible way.

The systems view of this creativity is articulated well by the creator of the system: “. . . the human and machine are considered as inseparable: two autonomous entities which, unavoidably, will influence each other, creating a unique meta-system made up of these two elements. The human and the machine establish a dialectics, a *talking through the other*, with no attempts of subordination, creating a performance which is the result of their cooperation, where, thus, the performer creates *together* with the machine.” (Sanfilippo, 2012).

Is this a computational creativity system? All of the sounds are coming from the system (but, the same is true for a piano). The human would not be able to make the work without the machine (but, the same is true of an artist without a pencil or whatever). Nonetheless, the computer/electronic system seems to be playing a stronger creative role in this interaction than that. Perhaps part of this is that the human is sometimes reacting to the outputs of the computer system as much as they are trying to shape it.

Contexts and Background Knowledge

In the list towards the beginning of the paper, we identified the background knowledge of the creator, and the context in which they were working in, as other things that could form part of what is responsible for a particular creative action or outcome.

Where is this background knowledge in a computational creativity system? In many cases it has been included as a part of the basic architecture of the system: for example, in Cohen's AARON system (Cohen, 1995), its figurative works are generated from parameterised algorithms that describe the basic figurative structures that are used to create the work. Other work draws on internet search algorithms as a way of accessing a background of knowledge (Johnson, 2013).

Can a way of accessing information, enabled by a technology, become part of the creative responsibility that a computer system provides to a creative activity or outcome? To what extent does the choice to use a complex, unpredictable computational technique in creating a work of art mean that that artwork has had a creative contribution from the computational system?

Let us consider a specific example. The image in Figure 1 is created by using the well-known Google image search functionality to search for images related to the word "secure" (filtered for images of a certain colour palette). If I choose to exhibit this as an artwork, where does the responsibility for the creative decisions sit? With me, alone? But I have hardly done anything! With the Google information retrieval system? With the people who have provided images for the system?

One role that computer systems—not just individual computers, but networked collections of computers with an associated infrastructure of information gathering and information retrieval—might play is to facilitate whole new areas of creativity. For example, the existence of vast online collections of images, together with technology of evergrowing sophistication to search and group such images by their meanings, facilitates a way of creating artworks that we might describe as *semantic mass*, where large collections of related information are gathered together and displayed.

Consider an example such as Jennifer Mills's work *What's in a Name?* (Figure 2), consisting of a large number of postcard-size paintings, each of which represents a person with the name "Jennifer Mills", gained from a search on Facebook. Is this work an artist's reflection on the ready ability to track down all of these people using the computer system, or is this a piece of collaborative creativity between the artist and that system? Even if it is not, does the system bear any "responsibility" for the artwork—any more than the paintbrush used to create the work?

Manovich (2002) has made related observations, that a technology can, by facilitating a change in the speed or scale of a process, create something which an observer might see as a genuinely new system. This can be seen by contrasting the Mills piece with comedian Dave Gorman's pre-web-search project *Are You Dave Gorman?*, where he tracked down a large number of people with the same name as him (Gorman, 2001). The Gorman project is focused on the labour of making the connections; the Mills piece on its effortlessness.

A number of artists and musicians have chosen to deliberately divest themselves of the responsibility of making creative choices in their art. Perhaps the best known of these is

John Cage, who created musical/theatrical works based on chance processes or on transcoding (Manovich, 2002) non-musical objects. An example of the latter is *Atlas Eclipticalis*, where star-charts were transcribed onto music staff paper, with stars representing notes, and the resulting music performed. By refusing the composer's traditional "responsibility" to decide (at a detailed level) where the notes go on a page, where has "responsibility" for the artwork gone? Perhaps an argument can be made that the responsibility has been abstracted to a higher level—that the details of the notes "don't matter", but the choice of star maps, rather than any any other printed material, is where the creator has chosen to vest his responsibility. A version of this argument has been made by Xenakis (1992), presenting a form of music in which the composer manipulates large-scale parameters of generative algorithms, rather than details.

There is a connection too, to the ideas of Goldsmith (2011), who has discussed the idea of "ostensive creativity", i.e. a means of being creative by "pointing at" material in the world, or organising it in a way that makes us see it afresh. Internet search based art can be seen as a form of this. But, who is doing the pointing?

Again, we are drawn back to a system view of the idea of creative responsibility. All of these components have some bearing on the final creative activity, and it is their interactions that lead to creativity happening, rather than one playing a responsible role and the others a supporting role.

Conclusions

At first it seems easy to distinguish between a system that is a tool that can be used in the aid of creative action, and one that takes on the responsibility for the creative act itself. However, when we look at complex, resistant artistic materials, systems containing complex interactions between humans and computers, and the kinds of human creativity and relational creativity that depend irreducibly on computers or networks of computers, then the distinction between a responsible creative agent, a creativity support system, and the more complex kind of tool become rather blurred.

It is easy to understand why the idea of responsibility finds its way into a definition of computational creativity. There is always a sneaky suspicion in a system involving interaction between humans and computers that all of the creativity is "coming from" the human (even when that human demonstrates surprise at the output from the system!). There is also the desire to distinguish creative systems from "mere tools". It is fairly clear that this can be done, up to a point, but the point at which tools slip over from being passive to being an active player in the creative process is a rather vague one.

Indeed, it is precisely because computers can be used to build complex, interactive, indeterminate systems that this distinction starts to become more problematic. Indeed, it is perhaps naive to assume that even in traditional non-computational artistic and musical creativity that a simple distinction can be drawn between individuals responsible for their creative action and the tools and concepts that they make use of. After all, reams of pages are written that attempt to explain why a particular artistic action was done by

contextualising it in the political, economic and social situation in which it is created.

One alternative approach would be to apply a “systems thinking” approach (Churchman, 1968) to this question. This approach would argue that it is futile to try and assign a particular component of the “art creating system” the definitive responsibility for producing the art work. Instead, there is a complex system of interacting agents and properties that lead to the work being realised (or not!) in the form that it ends up. By doing this, we are not throwing our hands in the air and saying that nothing can be said about how the work is produced. Instead, we are arguing that there is a complex system of interactions which in itself needs to be studied. Indeed, Csikszentmihalyi (1988) has explored a similar approach to explaining human creativity.

Perhaps we can modify the Colton/Wiggins definition in the following way: “the philosophy, science and engineering of computational systems which, by playing a role in an interactive system, contribute to that system producing behaviours that unbiased observers would deem to be creative”. Note that the systems have to “play a role” in the system; this opens up the possibility of many different possible roles.

This would seem to bring many activities that are currently seen as part of computational creativity squarely into the definition. For example, Veale (2011, 2013) has discussed the idea of creativity as a service, i.e. the provision of computational components that can be designed to be part of a larger creative system, glued together using web services frameworks.

The main point, however, is not to contribute to a pedantic (if sometimes enlightening) debate on definitions, but to shift the emphasis of computational creativity research. Rather than trying to identify the single actor in a complex, interactive system that is “responsible” for the creativity, instead we should recognise that this responsibility is diffuse and part of the behaviour of a complex human/computer system. That then leads onto much more interesting questions about how such systems gives rise to creativity, how components can be engineered for such systems, and how interactions in such systems can be managed, rather than searching for the single romantic hero who is the fount of all creativity in the system.

References

- Churchman, C. W. 1968. *The Systems Approach*. New York: Dell Publishing Co.
- Cohen, H. 1995. The further adventures of AARON, painter. *Stanford Electronic Humanities Review* 4(2).
- Colton, S., et al. 2009. Computational creativity: Coming of age. *AI Magazine* 30(3):11–14.
- Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier? In De Raedt, L., et al., eds., *ECAI2012 (Proceedings of the European Conference on Artificial Intelligence)*, 21–26. IOS Press.
- Csikszentmihalyi, M. 1988. Society, culture, and person: A systems view of creativity. In Sternberg, R., ed., *Nature of Creativity: Contemporary Psychological Perspectives*, 325–329. Cambridge: Cambridge University Press.
- Goldsmith, K. 2011. *Uncreative Writing: Managing Language in the Digital Age*. Columbia University Press.
- Gorman, D. 2001. *Are You Dave Gorman?* Ebury Press.
- Johnson, C. G. 2012. The creative computer as romantic hero? computational creativity systems and creative personæ. In Mayer, M. L., et al., eds., *Proceedings of the International Conference on Computational Creativity*, 57–61.
- Johnson, C. G. 2013. Artistic and musical application of internet search technologies: Prospects and a case study. *Digital Creativity* 24(4):342–266.
- Koza, J. R. 2010. Human-competitive results produced by genetic programming. *Genetic Programming and Evolvable Machines* 11:251–284.
- Lewis, G. E. 2000. Too many notes: Complexity and culture in voyager. *Leonardo Music Journal* 10:33–39.
- Manovich, L. 2002. *The Language of New Media*. MIT Press.
- McGregor, S. 2014. Considering the law as an evaluative mechanism for computational creativity. In al Rifaie, M. M.; Gow, J.; and McGregor, S., eds., *Proceedings of the Symposium on Computational Creativity at the 50th AISB Convention*. Available at <http://www.aisb50.org/>.
- Paine, G. 2002. Interactivity, where to from here? *Organised Sound* 7(3):295–304.
- Ritchie, G. D. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17(1):67–99.
- Rowe, R. 1993. *Interactive Music Systems: Machine Listening and Composing*. MIT Press.
- Sanfilippo, D. 2012. Lies (distance/incidence) 1.0: a human-machine interaction performance. In Polotti, P., et al., eds., *Proceedings of the 19th Colloquium on Music Informatics*, 198–199.
- Veale, T. 2011. Creative language retrieval: a robust hybrid of information retrieval and linguistic creativity. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 278–287.
- Veale, T. 2013. A service-oriented architecture for computational creativity. *Journal of Computing Science and Engineering* 7(3):159–167.
- Xenakis, I. 1992. *Formalised Music*. Pendragon Press. Second edition in English translation, first edition in French 1963.