The Creative Computer as Romantic Hero? Computational Creativity Systems and Creative Personæ

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

Abstract

A popular definition of computational creativity is that it consists in behaviour that would be regarded as creative if performed by humans. This raises the question of *which* humans, as there are many different styles of human creative behaviour. This paper unpacks a number of ways in which human artistic creativity can be characterised, compares them with the kinds of creative actions found in computational creativity, and explores some aspects of human creativity that are underrepresented in computational creativity systems.

Introduction

Computational creativity (CC) has been pursued for decades, with intense activity in recent years. One of the most common definitions of CC is "building software that exhibits behavior that would be deemed creative in humans" (Colton and others 2009). This paper explores this definition.

Creative behaviour is not monolithic; there are many different ways in which to validly be a creative person. It might be asked what *kind* of creative behaviour in humans we want to compare computer systems with. This paper is concerned with creative systems in artistic domains such as visual art, music and literature, not with so-called *everyday creativity*.

Creative Personæ

There are clear parallels between the idea of CC as exhibiting human-like behaviour and with Turing test-style definitions of intelligence. In both, the system is designated as acting in a creative/intelligent way if it can generate behaviour or product that would require, creative/intelligent action in humans to produce it, that a human observer would recognise.

Creative acts in artistic domains are not actions that most people will carry out regularly. So, care must be taken in selecting which humans are taken as exemplars. In traditional Turing style AI tests, the exemplar human is a member of the general population. In these tasks we will have to assume that the exemplar has some specific skills and knowledge.

One approach is to choose beginners in the domain as exemplars. This fits with the approach to CC system development that sees the development of systems that can do beginner tasks as the first step towards the development of more sophisticated systems. Another approach is to build systems to be compared with mature creative work, where the exemplars are mature artists. There is a danger with any of these exemplar-based systems (Pease and Colton 2011) that they encourage pastiche; but, if evaluators are primed sufficiently, perhaps this can be avoided.

McCormack (2005) notes that CC algorithms will be valued when they "produce art recognized by humans for its artistic contribution (as opposed to any purely technical fetish or fascination)". This seems, reasonable; yet, it might just be a temporary state. Once we have got beyond the point at which the products/activities of computational artistic creative systems are acknowledged as valid artworks, we might become interested in "biographical" aspects of them, and produce works that reflect on the origins of the work without this seeming like "technical fetish". This reflection on origins might become part of the depth of the works.

An important point is that not all creative people are creative in the same way. This paper will consider a number of dimensions of what will be termed *creative persona space*: an informally defined space representing broad attitudes/approaches. This paper considers three dimensions: the social vs. individualist dimension; the importance (or not) of ongoing tradition, development and "craft skills"; ideas of new and old media and the way in which technology is used in the artistic production itself.

Ritchie (2007) uses the *inspiring set* in evaluating CC: a set of human *works* as exemplars of what a successful CC system would generate. The intention here is similar, but with regard to the *creator* rather than the works: for a particular CC system, can exemplars of the creator that is being represented by the system be given?

Dimension 1: Socially Embedded vs. Individualistic

One dimension of difference between creative artists is between those that work as individuals and those that create work in a socially embedded context. No creative artist is entirely divorced from social context, but we focus on those that work directly with others in collaborative creation.

This is rare in the literary arts and uncommon in the visual arts; occasionally small groups will work consistently together over a long period of time (e.g. artistic duos such as Thomson & Craighead, the Chapman brothers and Cardiff & Miller), but literature, visual art and theatre/film writing are dominated by individual creators (comedy writing—

especially for TV—is a notable exception, as are the works of groups such as the *Dogme 95* filmmaking collective). In music, collaboration is more common. This reaches its peak in performance styles such as free improvisation, where a number of performers work together to collaboratively create a work without a preformed plan or an idea of leadership or direction. In some more commercial creative domains, such as advertising, group creative work is standard.

Most work in CC focuses on the individualistic concept of creativity: writing of stories, creation of jokes, composition of melodies, creation of pictures. There has been some interest (Cook and Colton 2011) in mapping out the various contributors to creativity in CC; this is explored futher below in the discussion of computer as medium. Whilst CC systems might create interactive works—for example in game level design (Togelius and others 2010)—it is rare that the system *continues* to be creatively active during interaction.

There are a number of potential reasons for this focus on the individual. These are readily criticisable and there is no reason to believe that all of them are believed by all practitioners, but listing them gives us an initial scoping out of the potential reasons:

- The work in CC is coming out of an AI tradition, which has focused on the idea of the simulation of the individual mind interacting with a task (though multi-agent systems are a counterexample).
- From an artistic perspective, there is a tradition of the "lone genius" in the romanticist tradition in art (Lovejoy 1948), which sees the role of the artist as developing their own authentic and original voice. This idea of the artist as *romantic hero* rejects the idea of collaboration and the development of an ongoing tradition, instead of which the great artist is seen as creating an individual body of work expressing their own personal world-view.
- Creativity might be seen as happening because of various interacting processes within the mind (see e.g. Koestler (1964)). However, people have a curious reluctance to admit hierarchical models of interacting networks, both in intelligence and creativity: creativity/intelligence might be seen as a product of interactions, but it is tempting to contain those interactions to one *level* in a complex system. It is difficult to conceive of a system where creativity is a product of interacting systems within the mind and *also* a network of interacting minds.
- There is a desire to be able to pin down exactly where the creativity is coming from. If a system is embedded in a complex social system with both human and computer agents, it is harder to point to a specific creative act by the computers. An easy criticism of such systems is that all of the creativity is coming from the human agents, and the seeming creativity of the computer agents is replicating, decorating or making trivial responses to or elaborations of the creative acts of the human agents.
- Individual creativity might be seen as the first stage in the development of more sophisticated interactive creative systems; therefore, until CC systems have demonstrated individual creativity, there is no point in tackling the "more complicated" task of group creativity.
- In practice, most creative work is individual; group creativity, in the arts, is confined to specialised areas. Therefore, CC systems are just emulating the world.

There are a number of such collaborative systems that have been created. Consider *Voyager* (Lewis 2000), where heuristics interact within a listening/responding musical system that improvises alongside human musicians; and Sanfilippo's LIES (Sanfilippo 2012), where sound processing systems are connected, the parameters of these interactions being adjusted by the user. Is this a CC system? It is easy to say that the creativity is coming from the user in the form of the parameter changes; but, those might be provoked by sounds from the system, what Blackwell and Young call "strong interactivity", which "depends on instigation and surprise as well as response." (Blackwell and Young 2005).

What might a CC system that was designed to work in a free, collaborative environment look like? Take musical improvisation as an example. One source of inspiration might be the broad guidelines that are given to beginners making a start in improvisation; whilst improvisation might be "free", it is not "anything goes" and there is a strong, often unarticulated, tradition about what is acceptable behaviour in such performances. This needs to be learned by participation and reflection, but guidelines can help to guide beginners so that they are not floundering totally.

Consider, the three guidelines put forward by Dave Smith: Listen, Don't waste sounds, and Develop a sense of social responsibility. How might a CC system attempt to work with these? "Listen." is both trivial (we need to have some means of getting input from the other performers) but, of course, is actually a very deep and complex guideline, especially considering that listening is an active process. Clearly, "listen" means that improvisation should take account of that listening. However, account-taking can easily become trivial, and bejust imitation; learning to *develop* material, and links between different heard sounds, is important to produce depth.

How might a CC system "listen" in this deeper sense? One characteristic of listening is that people frame listening with regard to what has been heard in the past, making subtle distinctions between some objectively very similar sounds, and grouping other sounds together that are objectively different (e.g. (Goto 1971)). One way to do this would be to run a system for a long time, and accumulate a set of listenings. This runs into the problem (Bown and McCormack 2010) of getting people to interact over a long period of time with a "naive" system that isn't providing engaging feedback.

An alternative would be to take inspiration from the idea of an *adult learner* who is new to free improvisation but already has a body of sonic knowledge from listening, speaking and playing an instrument. For example, a system might match listened phrases to a corpus of sonic information—a set of melodies, or a set of nature sounds like birdsong, or an artificially generated set such as sound contours derived from spoken text. Developments in audio information retrieval might mean that such information could be gained from web search. The system could then base responses on these matches, which would bring in a broader, allusive set of responses than those that just work with varying the input based on just that input alone.

Consider the second guideline: "Don't waste sounds.". Again, there is a naive interpretation of this: don't play all of the time. This is simplistic, but could provide the basis for an implementation; indeed, this is how many beginning improvisers might, with some success, interpret it.

However, it has more depth. "Waste" could include aspects of listening: don't waste the other sounds that are going on in the environment, whether by ruining them with your own sounds or by failing to exploit the sounds that have potential. There is a suggestion of depth of contribution: providing a contribution that is constantly striving for more depth and development and eschewing cliché,

Again, simple versions of these could be implemented in a CC system, particularly when combined with the deeper listening ideas above. For example, the agent could measure the complexity of the current activity and hold back in over-complex situations, the agent could hold a medium-term memory of what material is being worked on and wait until a particular piece of material is being "played out" before introducing new material, and so on.

Turning to the final point, how could a computational agent "develop a sense of social responsibility"? The "society" in question clearly includes the fellow players; also perhaps, the audience. One aspect of being socially responsible is recognising when behaviour is regarded by others as gauche or inappropriate. Could a computer agent measure this by the inference of affective states from fellow players (Picard 1997)? Clumsily, this could be done by direct feedback from the players to the agent. Another way in which humans learn social cues is by observing social interactions amongst others; could a creative agent build a model of such interactions within the group of improvisers and then use this model to develop its own interactions? Or, more simply, by analyse whether the material it performs is taken up by fellow players and use that as a proxy for the appropriateness of certain kinds of interaction?

Finally, how could such a system be evaluated? Evaluation of "individualist" creative systems is typically done via the impression that the system makes on the audience. However, a collaborative system could also be evaluated by the fellow players, and such players might use different criteria to that used by (or, indeed, perceptible by) an audience. For example, a creative improvisation system might be rated highly by fellow (human) improvisers if it provides a contextually-sensitive way of provoking the other improvisers to produce more engaging music.

Dimension 2: Tradition-Classicism vs. Romantic Individualism

Another aspect to the notion of the "romantic hero" is the contrast between the artist working in isolation as an individual genius and the artist working in a tradition. This is somewhat different to the above discussion: this section considers the contrast between the artist who pursues their own "individual voice" against the idea of one who is concerned with developing out from and contributing to the development of an ongoing tradition. By contrast, consider a *classicist* view which sees originality happening via the gradual development of a tradition, constantly underpinned by ideals of balance and proportion.

The idea of *romantic originality* was a significant shift in concept of artistic development within the European mainstream tradition, contrasting with earlier traditions about artistic creation being about skillful execution within a style, including reuse and redevelopment of material. The romantic artist creates work that reflects their own, individual, often troubled, engagement with the world (Butler 1981).

The relationship of CC to this dimension is complex. The idea of creativity expressed by Boden's model of *transformational creativity* reflects the classicist notion of a gradually developing tradition; the space is, after all, *transformed*, not *rejected*. But, perhaps too much stall shouldn't be placed on this—after all, all artistic work builds to some extent on previous work, and the most individualist romantic hero uses tools developed by previous generations. Indeed, individual initiative eventually tips over into eccentricity; witness the reviews collected by Slonimsky (2000), or the reception of "outsider art" (Rhodes 2000).

In areas where there is no ground truth, radical novelty is difficult to evaluate. A seemly incomprehensible piece could be madness or genius. In Boden's terms, if a transformation consisted of taking one space and replacing it by another, how would works in this new space be evaluated? For a more sober transformation, evaluation can start with our existing ideas of evaluation and push them a little; in a completely new space there is no corresponding grounding.

The view along this axis therefore gives a contrast to the previous one. Typical CC systems represent individual creativity rather than social creativity; but, they represent individual creativity within a tradition rather than that of the radical outsider.

Dimension 3: Old vs. New Media

Another distinction that can be made in artistic creativity is between so-called *old media* and *new media*. Defining new media is complex and contestable. Manovich presents an initial, naive understanding of new media as those cultural objects that essentially involve the use "of a computer for distribution and exhibition." (Manovich 2001).

He goes on to describe ways in which digital technology has influenced the *process* of creating cultural works. One example is where speedup faciliates a difference in kind rather than just a difference in degree. E.g., real-time rendering of 3D scenes makes interactive games possible as well as improving the production process of traditional animation. Furthermore, computer technologies provoke artists into exploring new creative areas: for example, the notion of *transcoding* (Manovich 2001), i.e. the ready exchange of data between different media formats, makes us think about creating works in which different media streams are created from the same source material, whether in a supportive or disruptive manner.

Computational Creativity: Old or New Media?

Are the products of CC systems *old media* or *new media*? Perhaps counterintuitively, the vast majority of CC systems are in an old media tradition. Whilst the computer is essential in CC, the role it plays is as *creator*; in new media, the computer is essential in the work *as it is presented*.

The work on computational creation of stories (Gervás 2009), poetry (Manurung and others 2000) and jokes (Ritchie 2009) is clearly in this vein: the aim of the vast majority of such systems is to create a work that is presented as words-on-the-page; whether these words are presented on paper or on the computer screen (or read out

loud) is not part of their essence. There are a few examples of literary creativity systems that *are* clearly within the new media tradition: for example, *nn* by Montfort (2007) is an example of the generation of interactive fiction.

In the case of CC systems for music, the landscape is more mixed. Systems such as *Voyager* produce sequences of notes to be performed by a synthesis system (sounding like a traditional instrument) or by a mechanical instrument. However, there are a number of examples that demonstrate how a creative music computing system could work in a new media fashion. For example, Magnus's *evolutionary music concrète* experiments (Magnus 2006) and Sanfilippo's LIES system discussed above show how creative computer systems can create electronic music that is concerned with sound manipulation rather than the production of notes.

CC works in a real new media tradition are rare. There is little CC work producing internet art (Greene 2004) or multimedia works. There are few works that use computationally-based means of organising material such as transcoding, or dynamic creation of work from a database (Manovich 2001), or whose aesthetic is a computational one such as the *database aesthetics* discussed by Vesna (2007).

One example is *Dance Evolution* by Dubbin & Stanley (2010), which uses a computer game engine as the basis for a system whereby characters learn to dance in time to music. The characters are stock video-game images of soldiers; it is not clear whether this was simply a use of the resources that were readily available within the engine, or whether the choice was deliberate. Regardless, this unusual choice of avatar provides a provocative image, reminiscent of various artists attempts to subvert the video-game culture by, for example, the performance of street theatre within MMORPG environments (Greene 2004).

The Creative Networked Computer

Despite CC arising during the "Internet age", the typical CC system produces creative output in a fixed medium (e.g. words, pixels on a screen, MIDI notes) that ignores the networked context of the computer. CC systems that draw on allusions to the world beyond are rare. Most storytelling systems (Gervás 2009) produce stories about a fixed set of ideas and characters. Most of the CC systems for visual arts work in an abstract medium.

Where they *do* provide external reference, this has typically been provided by the system designer directly. One way is that the designer builds into the system some understanding of the external world: for example, whilst the details of how people are drawn in Cohen's *AARON* system (Cohen 1995; Boden 1990) are created by the system, the basic idea of a person-shape is part of the system design. The second way this that the designer might place processes within the system that generate allusions to the world beyond, for example in *ecologically* inspired systems such as those discussed by Bown & McCormack (2010) where the interactions between components of the system are inspired by the kinds of interactions found in natural ecological systems. This could well suggest aspects of the natural world to viewers of the system, even if the presentation is very abstracted; but, the decision to make this allusion is that of the system designer.

There are exceptions. Krzeczkowska et al. (2010) present a system where the initial source material is drawn from current news stories, and keywords extracted are used in web searches for pictures, which are used as the source material for the creation of a collage using Colton's *Painting Fool* system (Colton 2012).

An important part of much art (particularly visual and conceptual art) is connotation: the depth of the work comes from ideas that are suggested, triggering connections that remain under conscious awareness, or revealed via "ah-ha" moments where the link or allusion is suddenly revealed (the idea of CC systems *framing information* as part of their creativity (Colton and others 2011) captures some of this).

Computer as Creator; Computer as Medium

CC researchers might eschew working in new media because of the potential for confusion between two different roles for the computer. For a CC system in a computer-based medium, the computer is playing two roles; that of creator, and that of the medium in which the work is realised. In theory, there is no reason why such a system should not be successful. However, in terms of *evaluation*, the creative computer working in new media is harder to evaluate. There is a question of the role that the "creative" computer played, versus the role played by the broader computational context.

McCormack (2005) has argued that to be taken seriously, CC systems need to create works that are not just examples of "technical fetish", but that are accepted in their own right. This may be why CC systems have avoided new media works: many new media works play with the idea of their technological groundedness in a self-referential and essential way, and CC system creators avoid building systems that work in such media to avoid accusations of "mere" technical obsession. But, awareness of origins need not reflect an over-obsession with a trivial part of the work; indeed, the *lack* of any sense of biographical depth is one of the shallowing aspects of many CC-produced works.

One new direction would be to build CC systems that explore and celebrate what computers are good at. One inspiration for this comes from John Cage's account of his development:

> After I had been studying with him for two years, Schoenberg said, "In order to write music, you must have a feeling for harmony." I explained to him that I had no feeling for harmony. He then said that I would always encounter an obstacle, that it would be as though I came to a wall through which I could not pass. I said, "In that case I will devote my life to beating my head against that wall." (John Cage, *Lecture on Indeterminacy* (Cage 1973))

A CC system could adopt a similar attitude, to acknowledge that computers might not be capable of "passing off" certain aspects of human creativity (e.g. creating fluent natural language text) but that certain new forms of creativity are facilitated that draw on the specific capabilities of computers and the computational infrastructure.

Let us consider two recent works. Kessel's 2011 *Photography In Abundance* where a day's uploads to *Flickr* were printed out and placed in a gallery—hundreds of thousands of photos. Such a work could not have been achieved without computational infrastructure; yet, is this *computational creativity*? Presumably not, as the creative decision about

what to do with the computational capability of Flickr was made by the human artist. Yet, the computational infrastructure seems to have played a stronger role here than that of medium—the contribution of the computer to this work is greater than the contribution of paint to a painting. Perhaps we need to get used to the idea of artists working in a richer medium, which "pushes back" in terms of creative contribution much more than a conventional medium.

The second example is my own 2011 piece *Blank*: nine printed panels containing image search hits for "blank". Most of the images are of empty objects: map outlines, blank signs, empty music paper. One image in the seventh panel consists of a photo of a number of gun cartridges—referencing "blanks" as a cartridge without a bullet.

Had this work been produced unaided by a human artist, people might ascribe it creative depth. Having set up the expectation of empty, neutral images, there is this single flash of violence in the middle of one panel, causing us to reinterpret the meaning of the emptyness. Of course, there is no *intention* to create this as such; it is "mere" serendipity. Yet, it is deeper than the readerly (Barthes 1970) interpretation of a purely aleatoric work; the images were "chosen" by a process that has a huge amount of infrastructure behind it. Is this any shallower than some purely human-produced work which has "come to mind" as the end result of the memory structures produced by the artist's life experiences?

References

Barthes, R. 1970. *S/Z*. Editions du Seuil.

Blackwell, T., and Young, M. 2005. Live algorithms. *AISB Quarterly* 122:7–9.

Boden, M. A. 1990. *The Creative Mind, Myths and Mechanisms*. Weidenfeld.

Bown, O., and McCormack, J. 2010. Taming nature: Tapping the creative potential of ecosystem models in the arts. *Digital Creativity* 21(4):215–231.

Butler, M. 1981. *Romantics, Rebels and Reactionaries: English Literature and its Background 1760-1830*. Oxford University Press.

Cage, J. 1973. *Silence*. Marion Boyars.

Cohen, H. 1995. The further exploits of AARON, painter. *Stanford Humanities Review* 4(2).

Colton, S., et al. 2009. Computational creativity: Coming of age. *AI Magazine* 30(3):11–14.

Colton, S., et al. 2011. Computational creativity theory: The face and idea models. In *Proceedings of the 2011 International Conference on Computational Creativity*.

Colton, S. 2012. The painting fool. http://www.thepaintingfool.com/ (visited April 2012).

Cook, M., and Colton, S. 2011. Automated collage generation—with more intent. In *Proceedings of the 2011 International Conference on Computational Creativity*.

Dubbin, G., and Stanley, K. O. 2010. Learning to dance through interactive evolution. In Di Chio, C., et al., eds., *Applications of Evolutionary Computation*, 331–340. Springer.

Gervás, P. 2009. Computational approaches to storytelling and creativity. *AI Magazine* 30(3):49–62.

Goto, H. 1971. Auditory perception by normal japanese adults of the sounds "l" and "r". *Neuropsychologia* 9(3):317–323.

Greene, R. 2004. *Intenet Art*. Thames and Hudon.

Koestler, A. 1964. *The Act of Creation*. Macmillan.

Krzeczkowska, A., et al. 2010. Automated collage generation—with intent. In Ventura, D., et al., eds., *Proceedings of the International Conference on Computational Creativity*, 36–40.

Lewis, G. E. 2000. Too many notes: Computers, complexity and culture in *Voyager*. *Leonardo Music Journal* 10:33–39.

Lovejoy, A. 1948. On the discrimination of romanticisms. In *Essays in the History of Ideas*. Johns Hopkins University Press.

Magnus, C. 2006. Evolutionary musique concrète. In Rothlauf, F., et al., eds., *Applications of Evolutionary Computing*, volume 3907 of *Lecture Notes in Computer Science*, 688–695. Springer Berlin / Heidelberg.

Manovich, L. 2001. *The Language of New Media*. MIT Press.

Manurung, H. M., et al. 2000. Towards a computational model of poetry generation. In *Proceedings of AISB Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science*, 79–86.

McCormack, J. 2005. Open problems in evolutionary music and art. In Rothlauf, F., et al., eds., *Applications of Evolutionary Computing*, volume 3449 of *Lecture Notes in Computer Science*, 428–436. Springer Berlin / Heidelberg.

Montfort, N. 2007. Ordering events in interactive fiction narratives. In Magerko, B. S., and Reidl, M. O., eds., *Intelligent Narrative Technologies: Papers from the 2007 AAAI Fall Symposium*, 87–94. AAAI Press.

Pease, A., and Colton, S. 2011. On impact and evaluation in computational creativity: A discussion of the turing test and an alternative proposal. In *Proceedings of the AISB symposium on AI and Philosophy*.

Picard, R. 1997. *Affective Computing*. MIT Press.

Rhodes, C. 2000. *Outsider Art: Spontaneous Alternatives*. Thames and Hudson.

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

Ritchie, G. 2009. Can computers create humor? *AI Magazine* 30(3):71–81.

Sanfilippo, D. 2012. Dario Sanfilippo. http://dariosanfilippo.tumblr.com/ (visited April 2012).

Slonimsky, N. 2000. *Lexicon of Musical Invective*. W.W. Norton and Co.

Togelius, J., et al. 2010. Search-based procedural content generation. In Di Chio, C., et al., eds., *Applications of Evolutionary Computation*. Springer. 141–150.

Vesna, V., ed. 2007. *Database Aesthetics: Art in the Age of Information Overflow*. University of Minnesota Press.