The Image in the Age of Computer Agency

[Draft of text – submitted for the Proceedings of the Impact 7 Conference, Monash University, September 2011]

Gordon Monro
Monash University, Australia
gordon@gommog.com

Abstract
Walter Benjamin identified the loss of the 'aura' associated with the uniqueness of an artwork that occurred when distribution of mechanical reproductions became prevalent. In this paper I discuss the loss of another aura, that of the authorial agency of the artist, brought about by handing over part of the creative work to a computer program. Here I am not concerned with collages, appropriation and the like, which can be seen as distributed authorship, but with the situation where part of the authorship can reasonably be assigned to a computational process that makes autonomous aesthetic judgements. Examples are discussed in areas ranging from everyday experience with digital media and technology to experiments on the frontiers of art.

Introduction

In 2006 a 45-metre long image entitled *Bloom* was displayed on a billboard at the Queensland University of Technology. The image was by Australian artist Jon McCormack. McCormack didn't print the final image, but he didn't create all the forms in the image either. Rather, the forms were evolved by a computer program written by McCormack (2006).

The history of printmaking is bound up with technical advances. The original advance, the ability to produce multiple copies of the same image, led to the loss of the 'aura' of a unique work of art, as identified by Walter Benjamin (1973). Now further technical advances have led to creation of works like *Bloom*, where decisions with aesthetic aspects are being made, not by a human artist, but by a computer program. A second 'aura', that of the authorial agency of the artist, is being eroded by technical developments.

A Double Intervention

In his book on Australian printmaking in the 1990s Sasha Grishin includes some prints made with laser printers. In this context he remarks that the use of computers and laser printers 'reasserted printmaking's traditional role of being the art form that which employs the most recent technology' (Grishin 1997, 11). But the computer brings with it more than just new machinery for printmaking. This is seen in a small handbook for artists for the Canon CLC-1 laser copier (Baraki and Smith 1991). The copier, which pre-dated Adobe's Photoshop, had some digital image-editing capabilities; in particular it could apply different colour transformations to different parts of the same image. Baraki and Smith go so far as to refer to the 'electronic plate', that is the instructions for the copier to process the scanned-in image.

Grishin considers that 'probably the most distinguishing feature which separates printmaking from the other major [visual] art forms lies in the processes involved in the production of a print ... the appearance of the finished print will often be dependent on unpredictable effects of the technical process' (Grishin 1997, 10). With the Canon copier the technical process is split into two essentially different parts: the computational manipulation of a digitally stored image and the physical creation of the image by the depositing of toner on paper. These two technical processes constitute a double intervention of technology into the image-making process. They also illustrate what Druckrey (1996,13) has well described as 'the evolution from mechanical to computational culture'. The key word here is 'computational', not 'electronic' or 'digital', as what is most characteristic of the computer is its ability to carry out complex calculations autonomously.

This double intervention has of course become a fact of life, and in most image-making contexts the computational intervention now overshadows the mechanical intervention. Thus a recent textbook (Dabner, Calvert et al. 2010) has the subtitle 'A foundation course for graphic designers working in print, moving image and digital media' and devotes only part of one chapter to the physical processes of printing. It is not surprising that artists like McCormack typically work with video and real-time interactive displays as well as prints.

Computational Opacity

There is a direct connection between the inscription of a line on a plate and the mark on the resulting print, even though the mechanical process of the printing press has intervened. With the computer, the connections between the inputs to a computational process and the outputs of the process can be obscure; indeed in everyday imaging the computational process is typically opaque to the user.

I have a small digital point-and-shoot camera. Such cameras come with several 'scenes', portrait, night, and so on. If I select one of these scenes the camera applies settings that alter the appearance of the captured image. But, beyond the 'scenes', the default setting of my camera is 'intelligent auto', where the camera itself chooses what scene to apply, based on properties of the image it has on its sensor. The camera is making an automated aesthetic judgement, based on principles that are opaque to me.

From image capture I turn to post-capture image processing. There are a number of touch-up tools in Adobe Photoshop, ranging from a straightforward copy-and-paste operation to the 'healing' tool, which blends a 'source' area and a 'target' area in a way that takes into account what is already present in the target area; the mode of operation of the tool is not obvious. When it was introduced its description promised to allow users to 'effortlessly remove dust, scratches, blemishes, wrinkles, and other flaws' (Pasini 2002). However disturbing the categorisation of wrinkles as flaws may be, for this paper the key word is 'effortlessly', indicating that the Photoshop user does not need to understand how the tool works.

The 'healing' tool has recently been upstaged by Photoshop's 'content-aware fill', introduced in 2010. This allows a user to select an unwanted part of an image, for example a tree or a person, and remove it; the program automatically synthesises background from the surrounding areas, and does it remarkably well, with very little input on the user's part. How this tool works is completely opaque to anyone without considerable expert knowledge, and indeed Adobe uses the work 'magical' on its website (Adobe Systems Incorporated 2011). Work on 'texture synthesis' and 'image completion' has been underway for some time in academic institutions (Wei et al. 2009; Komodakis and Tziritas 2006), but it has now entered the commercial arena with a vengeance.

One can distinguish three levels of computational opacity, where the operations of a computer are opaque to the user. The first is exemplified in such operations as drawing a straight line in a program like Adobe Illustrator. What to a human are natural operations on graphic images often do not translate well into the realm of pixels in a digital computer. Drawing a straight line is a case in point; see for example the discussion of Bresenham's algorithm and related methods in (Foley et al. 1990, 74–78). But since the convoluted digital process corresponds to a simple pencil-and-paper operation the user does not give it any thought. The second level of opacity occurs when the user is aware that an operation depends on principles that he or she does not understand, as with Photoshop's 'content-aware fill'. The most interesting is the third level of opacity, which occurs when an artist writes a computer program for artistic purposes, but the artist (now an artist-programmer) is not able to work out in detail how the program produces the results it does, even though he or she wrote the program. This situation is further discussed below.

Giving a Computer Program Aesthetic Knowledge

Image-manipulation tools such as those in Photoshop mentioned above take into account some features of the human perceptual system, but they do not make use of traditional compositional rules employed by artists and photographers such as the rule of thirds. Here I discuss three attempts to provide a computer program with the sort of aesthetic knowledge routinely employed by human artists (or on one case human composers).

Probably the best-known such attempt is the very long-running program *AARON* by Harold Cohen, who was a well-known painter before he began working with computers (McCorduck 1991). *AARON* began in the early 1970s as a drawing program that controlled a device holding a pen; from the outset it had aesthetic concepts built into it. Cohen states: 'The earliest versions of *AARON* could do very little more than to distinguish between figure and ground, closed forms and open forms, and to perform various simple manipulations on those structures' (Cohen 1995). Cohen goes on to say 'I was beginning to face the fact that the human cognitive system develops in the real world, not in the vacuum where *AARON* lived', which led Cohen to add information to the program on how to draw people and trees, and to use rules concerning perspective and three-dimensional placement (though the result was still a drawing on paper). *AARON* thus entered into a representational phase. Subsequently Cohen provided *AARON* with information about colour, and it entered into a colourful semi-abstract phase producing images that 'teeter on the edge between abstraction and representation' (Bernard Jacobson Gallery 2008). In each of these developments Cohen is incorporating human-developed aesthetic concepts into the workings of the program.

The second attempt I consider is *The Painting Fool* by Simon Colton and co-workers (Colton et al. 2008). They state: 'We are currently building an automated painter (called *The Painting Fool*) which we hope will eventually be accepted as a creative artist in its own right.' *The Painting Fool* is a 'non-photorealistic' program that produces digital images that look like drawings or paintings. Other programs exist that render a photograph in a way that makes it resemble a painting, but *The Painting Fool* aims to go well beyond this. In a striking demonstration, the program was coupled with another program that attempted to discern the emotion expressed in a short video of a face. *The Painting Fool* then rendered the face in a manner appropriate to the detected emotion, for example using a subdued palette to express sadness or abrupt lines to express anger. In this context, Colton has introduced what he calls the 'creative tripod', consisting of skill, appreciation and imagination; he argues that all three must be present for a perception of creativity. Colton is working towards giving his program all three attributes (Colton 2008).

Finally I consider an attempt to give a program aesthetic knowledge, not in the domain of images, but in the domain of music. A digital computer is much more adapted to work with music, or at least Western notated music, than it is to work with images: notated music employs a discrete set of pitches (pretty much the 88 notes of the piano keyboard); similarly the note durations and the degrees of loudness of the notes are limited to finite lists. Thus uses of the computer in music may well foreshadow analogous developments in image-making, and it is not surprising that the handing over of aesthetic judgement to computer programs has been taken further in music than in visual art. The most striking work in this area is by the composer and musicologist David Cope (2001). Cope created a computer program called *EMI* (standing for 'Experiments in Musical Intelligence'). To compose a work, the program is primed with a database of pieces, and then recombines elements from them to create a new piece, using sophisticated analyses created by the computer. Cope used the program to create a new Mozart symphony, new works in styles of Chopin and Rachmaninov, and 5,000 Bach chorales. In some cases the results were sufficiently good that experts could not tell the difference between genuine works by the relevant composer and EMI's compositions: Kala Pierson, a graduate student at the Eastman School of Music, one of the top American schools, reported that at a performance at Eastman a majority of composition and music theory people picked wrongly when asked to choose between a 'Chopin' mazurka by *EMI* and one actually by Chopin (Cope 2001, 66–67).

Artificial Life Art

The context for Jon McCormack's *Bloom* is quite different from that of programs which have some knowledge of traditional aesthetic criteria; instead aspects of the operations of the natural world, such as growth and development of plants or the behaviour of insects, are taken as inspiration. Programs so inspired have been given the general name 'artificial life'. Aside from art they have been used for scientific purposes, to cast light on aspects of real biological life, and also to provide methods of solution for problems in engineering.

A major source of inspiration for artificial life art has been the notion of biological evolution. A pioneering work in this area was Karl Sims's (1991) *Genetic Images* (first exhibited in 1993). Here each image was generated by a small piece of computer code that acted as 'DNA'. A form of 'survival of the fittest' was applied by visitors to the exhibition, who could choose the images that survived to be 'parents' for the next generation. The images in *Genetic Images* were abstract, but Sims himself and other artists including McCormack developed works based on plant and animal forms; in some cases they also moved to getting the computer itself to carry out the 'survival of the fittest'. Mitchell Whitelaw's book *Metacreation* (Whitelaw 2004) surveys the area.

Artificial-life principles were applied in the movie *The Lord of the Rings: The two towers*. The orcs were animated in part by a computer program, but not one of a conventional sort. Each individual computer-generated orc was given a 'brain' with quite an elaborate set of rules. There was no central controller; each orc animated itself. The interactions of the orcs, with each individual following its own rules, produce very convincing results (New Line Productions 2007). In this instance the computer is generating *behaviour* (the on-screen movement of the orcs) rather than form or colour, illustrating the wide reach of the computational intervention.

The fairly complex rules used in the sort of programs considered here, which are often structured as consisting of many largely autonomous sub-programs, can produce results that are difficult for even the artist-programmer who wrote the program to understand. Although such programs usually employ randomness to some extent, the difficulty in understanding is not due so much to the randomness as to the complex and conditional interactions among the parts of the program. Thus Cope has said in discussing *EMI*: 'the complex interactions between the various rules governing [construction of bars approaching a cadence] make such decisions impossible to reverse-engineer' (Cope 2001, 159). That is to say, no one, not even the creator of the program, can determine from the results how the program reached them.

Are the Computers Taking Over, and Do We Care?

With the exception of Cope's work, the developments described above have been accepted with remarkable equanimity. I think that this is because the new tools and procedures fit into two long-standing existing contexts for intellectual and creative work.

The first such context is the use of artificial (non-biological) aids to thought, as discussed by philosopher Andy Clark in his book *Natural-Born Cyborgs* (Clark 2003). Clark contends that, certainly since the invention of writing, much of our thinking takes place outside our biological brains: we use sketches, notes, diagrams, diaries and other aids such as dictionaries and encyclopaedias. Until recently, if we wished to multiply two large numbers together we used pencil and paper. Now the traditional aids to thought have been augmented by calculators, spreadsheets, spell-checkers, drawing programs, smart-phones and of course the vast resources of the World Wide Web, made instantly available by search engines. Clark sees all this as quite natural, arguing that our brains are well suited to make use of such external aids to thought, and have the plasticity to adapt readily to new tools as they come along.

The second long-established context is that of the master-apprentice relationship, and more generally that of a perceived creative hierarchy. In a large enterprise such as the production of a feature film or a modern computer game such a hierarchy is inevitable, as the bewildering list of roles in film credits indicates, but the master-apprentice relationship dates back to antiquity. A computer program can be regarded as playing the role of an apprentice or an assistant; already in 1979 the artist and critic Andrew Forge had suggested that *AARON* be considered Cohen's apprentice rather than a tool Cohen used (McCorduck 1991, 82).

Thus, since we are already used to employing non-biological aids in intellectual work, and we are also used to the idea that creative tasks considered as low-level may be delegated, it is not a big stretch to delegate some apparently creative tasks to machines. Nonetheless the introduction of computers acting as creative agents does bring in new situations. I consider three such, which I view as occurring in three different levels of the master-apprentice relationship.

Firstly I consider the lowest-level computer apprentices. The development of computer programs to automate low-level tasks in creative contexts has led to the observation that some of the activity hitherto considered as creative, albeit low-level, may actually be achievable in a very simple-minded way. Cheap tricks, or less pejoratively simple mechanisms, such as the rule of thirds can prove surprisingly effective.

The second situation is where there is a high level of delegation to the computer apprentices. This arises in complex artificial-life artworks such as *Black Shoals* by Lise Autogena and Joshua Portway (Blackshoals.net 2004) or *Eden* by Jon McCormack (2004) that create what amounts to small-scale ecosystems in the computer. The artist moves from being a creator to a meta-creator (to adopt Whitelaw's term) who designs the system as a whole but is not able to predict in detail the outcomes of the interactions in the system. In an extreme case we may have the situation described by Domenico Quaranta:

> The author [of a computer-generated artwork], therefore, sets into motion a process which develops itself autonomously, and, often, in an unpredictable way, under an amazed gaze. We seem thus to deal not as much with an artist, considered in the way we usually do, but rather with a minor God, who activates a system and then watch [sic] it coming to life. (Quaranta 2006)

This is largely a new situation; the nearest analogy is perhaps with the role of a gardener who lays out a garden and plants trees and bushes, but cannot direct details of their growth (McCormack and Dorin 2001). Even in this meta-creative case the question of authorship has not been raised seriously. The credit goes to the master, not the apprentices: the designer of the system, or analogously the landscape artist, not the computer program or the trees and bushes.

The third situation is where the mechanical apprentice is seen to usurp the role of the human master. Of the various works discussed above, the only one that has elicited really disturbed or antagonistic reactions is the output of Cope's *EMI*. Douglas Hofstadter has commented at length on his reactions to *EMI*, and comes to some depressing conclusions, the last of which is 'The *human soul/mind* is a lot shallower than I had ever thought'. (Cope 2001, 80; emphasis in original). He goes on to say

> [This possibility] would be the ultimate affront to human dignity. It would be the realisation that all the 'computing power' that resides in the human brain ... can by bypassed with a handful of state-of-the art chips, and that all that is needed to produce the most powerful artistic outbursts of all time (and many more of equal power, if not greater) is a nanoscopic fraction thereof—and that it can all be accomplished, thank you very much, by an entity that knows nothing of knowing, seeing, hearing, tasting, living, dying, struggling, suffering, aging, yearning, singing, dancing, fighting, kissing, hoping, fearing, winning, losing, crying, laughing, loving, longing, or caring. (Cope 2001, 80)

A comparable situation has not yet been reached in image-making, but there are some who would embrace it. The authors of the recent 'Istanbul Manifesto' declare:

> Manual skill leads only to a senseless waste of time. The human artist is not a maker, but a creator. Art is a mind extension, a prosthetic, a machine that just waits to be triggered. The role of the artist is to push the ON button, giving rise to an autonomous product. (Moura et al. 2011)

They conclude: 'The great artist of tomorrow will not be human'.

Conclusions

I draw three conclusions. Firstly, the aura of artistic creation survives the assistance of any number of computational apprentices, even if their contributions are elaborate and sophisticated, and even if their modes of operation are totally opaque, provided that there is a presiding human master. Secondly, if a work is presented in the context of high art as the work of a computer program the absence of the aura of human creation can be profoundly disturbing. And thirdly, attempts to endow computers with creative agency can potentially teach us a lot about the nature of *human* creativity, not least by dispelling mystification.

Benjamin begins his essay with a quotation from Paul Valéry, including the following words:

> We must expect great innovations to transform the entire technique of the arts, thereby affecting artistic invention itself and perhaps even bringing about an amazing change in our very notion of art. (Benjamin 1973, 219)

The transition from mechanical to computational culture is bearing out Valéry's prediction in a way that neither Valéry nor Benjamin could have foreseen.

References

Adobe Systems Incorporated. 2011. 'Photoshop CS5 Extended: Features.' Accessed 21 May 2011. Available from: http://www.adobe.com/products/photoshopextended/features._sl_id-contentfilter_sl_features%2Bby%2Bversion_sl_new%2Bsince%2Bcs4.html.

Baraki, Bashir; Smith, Gill. 1991. *Creating Laser Fine-Art Prints*. [Caulfield East]: The Art School Press.

Benjamin, Walter. 1973. The Work of Art in the Age of Mechanical Reproduction. *Illuminations* (ed. Hannah Arendt): 219–253. London: Fontana. (First published 1936.)

Bernard Jacobson Gallery. 2008. 'Harold Cohen: Colour rules'. Accessed 21 May 2011. Available from: http://www.jacobsongallery.com/content.php?flash=yes&nav=pressdetails&ID=20.

Blackshoals.net. 2004. 'Black Shoals : Stock market planetarium'. Accessed 21 May 2011. Available from: http://www.blackshoals.net/.

Clark, Andy. 2003. *Natural-born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford University Press.

Cohen, Harold. 1995. 'The further exploits of AARON, painter'. *Stanford Humanities Review* 4 (2): 141–158.

Colton, Simon. 2008. 'Creativity versus the perception of creativity in computational systems'. Presented at the AAAI [Association for the Advancement of Artificial Intelligence] Spring Symposium on Creative Intelligent Systems, Stanford University, California.

Colton, Simon; Valstar, Michel F; Pantic, Maja. 2008. 'Emotionally aware automated portrait painting'. Proceedings of the 3rd International Conference on Digital Interactive Media in Entertainment and Arts: 304–311. New York: Association for Computing Machinery.

Cope, David. 2001. *Virtual Music: Computer Synthesis of Musical Style*. Cambridge, Massachusetts: MIT Press.

Dabner, David; Calvert, Sheena; Casey, Anoki. 2010. *Graphic Design School : A Foundation Course for Graphic Designers Working in Print, Moving Image and Digital Media*. London: Thames and Hudson.

Druckrey, Timothy, ed. 1996. *Electronic Culture: Technology and Visual Representation*. New York: Aperture Foundation.

Foley, James D; van Dam, Andries; Feiner, Steven K; Hughes, John F. 1990. *Computer Graphics: Principles and Practice*. 2nd ed. Reading, Massachusetts: Addison-Wesley.

Grishin, Sasha. 1997. *Australian Printmaking in the 1990s: Artist Printmakers: 1990–1995*. Sydney: Craftsman House.

Komodakis, Nikos; Tziritas, Georgios. 2006. 'Image completion using global optimization'. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), vol 1: 442–452.

McCorduck, Pamela. 1991. *Aaron's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*. New York: W.H. Freeman & Co.

McCormack, Jon. 2004. *Impossible Nature: The Art of Jon McCormack*. Melbourne: Australian Centre for the Moving Image.

McCormack, Jon. 2006. 'Bloom'. Accessed 21 May 2011. Available from: http://www.csse.monash.edu.au/~jonmc/projects/Bloom/Bloom/bloom_main_page.ht

ml.
McCormack, Jon; Dorin, Alan. 2001. ‘Art, emergence and the computational sublime’. Second Iteration: Second International Conference on Generative Systems in the Electronic Arts. Melbourne: Monash University. [CD-ROM]. (Available online at http://www.csse.monash.edu.au/~jonmc/research/Papers/art-2it.pdf; accessed 21 May 2011.)
Moura, Leonel; Pereira, Henrique Garcia; Rinaldo, Ken. 2011. ‘The Istanbul Manifesto’. Accessed 21 May 2011. Available from: http://www.leonelmoura.com/manifesto_istanbul.html.
New Line Productions. 2007. ‘Helm’s Deep: Massive brains’. Accessed 21 May 2011. Available from: http://lordoftherings.net/effects/ (click on ‘Enter Helm’s Deep’, then ‘Armies’ and then ‘Massive Brains’).
Pasini, Mike. 2002. ‘More than Magic: Photoshop 7.0 in Action’. *Imaging Resource*. Accessed 21 May 2011. Available from: http://www.imaging-resource.com/SOFT/PS7/PS7.HTM.
Quaranta, Domenico. 2006.‘Generative ars [sic]: Catalogue essay for the exhibition C.STEM: Art Electronic Systems and Software Art’. Turin: Teknemedia. (Available online at http://domenicoquaranta.com/2009/09/generative-ars/; accessed 21 May 2011.)
Sims, Karl. 1991. ‘Artificial evolution for computer graphics’. *Computer Graphics* 25 (4): 319–328.
Wei, Li-Yi; Lefebvre, Sylvain; Kwatra, Vivek; Turk, Greg. 2009. ‘State of the art in example-based texture synthesis’. *Eurographics 2009, State of the Art Report, EG-STAR*. Accessed 21 May 2011. Available from: http://www-sop.inria.fr/reves/Basilic/2009/WLKT09.
Whitelaw, Mitchell. 2004. *Metacreation: Art and Artificial Life*. Cambridge, Massachusetts: MIT Press.