# 1

# *An Extract From...* Facing the Future: Evolutionary Possibilities for Human-Machine Creativity

Jon McCormack

Centre for Electronic Media Art
Monash University
Clayton 3800
Victoria, Australia
jonmc@csse.monash.edu.au
www.csse.monash.edu.au/∼jonmc

This is an extract from the chapter:
McCormack, J.: Facing the Future: Evolutionary Possibilities for Human-Machine Creativity, in P. Machado and J. Romero (eds.) *The Art of Artificial Evolution: a handbook*, pp. 417-451, Springer, 2008.
To see the full chapter you'll need to purchase the book, which is available from Springer. See:
http://art-artificial-evolution.dei.uc.pt/

**Summary.** This chapter examines the possibilities and challenges that lie ahead for evolutionary music and art. Evolutionary computing methods have enabled new modes of creative expression in the art made by humans. One day, it may be possible for computers to make art autonomously. The idea of machines making art leads to the question: what do we mean by 'making art' and how do we recognise and acknowledge artistic creativity in general? Two broad categories of human-machine creativity are defined: firstly, machines that make art like, and for, humans; and secondly, machines that make 'art' that is recognised as creative and novel by other machines or agents. Both these categories are examined from an evolutionary computing perspective. Finding 'good' art involves searching a phase-space of possibilities beyond astronomical proportions, which makes evolutionary algorithms potentially suitable candidates. However, the problem of developing artistically creative programs is not simply a search problem. The multiple roles of interaction, environment, physics and physicality are examined in the context of generating aesthetic output. A number of 'open problems' are proposed as grand challenges of investigation for evolutionary music and art. For each problem, the impetus and background are discussed. The paper also looks at theoretical issues that might limit prospects for art made by machines, in particular the role of embodiment, physicality and morphological computation in agent-based and evolutionary models. Finally, the paper looks at artistic challenges for evolutionary music and art systems.

> In the last analysis all intelligibility and all intelligent behaviour must hark back to our sense of what we are, which is, necessarily, on pain of regress, something we can never explicitly know.
>
> — Dreyfus and Dreyfus [1]

## 1.1 Everyimage

People understand the expression 'finding a needle in a haystack' as indicating that a problem is very difficult because it involves searching through a lot of things (probably too many) to find what one is after. Evolutionary music and art might be described this way — as researchers, we're trying to find the aesthetically satisfying needle from the data haystack of computation and algorithm. As we shall see, this is an understatement. So let us begin with a simple thought experiment. How difficult it is to find the good art using computational representations and processes? How hard is it to find the *Mona Lisa*[1] from the set of all possible images?

We will restrict our art to be two-dimensional, pixel-based images (the standard way images are stored on a computer). Initially this seems like a relatively simple subset of what might constitute art. However, the following argument can be adapted to any form of digitally representable media, irrespective of type or resolution (including non-visual representations, such as music).

The discrete, pixel-based image is curiously deceptive in terms of its complexity. A modest 500 × 400 pixel image, for example, contains only 200,000 pixels and is easily stored and manipulated on any modern computer. However, the space of possible images that can be represented within those 200,000 pixels is, as we shall see, *Vast*.[2]

[1] Readers might be perplexed with my continual use of the *Mona Lisa* as an exemplar artwork in this chapter. I am not suggesting researchers try to evolve an image that is literally the *Mona Lisa*. I use the *Mona Lisa* for a variety of reasons, including a kind of postmodern irony. In the canons of Western Art, this image is broadly recognised as an exemplary example of fine art. It has been around for long enough to be reasonably sure it is not the product of a fad or distortion of what constitutes art, albeit in a classical, Western sense. It is widely known and instantly recognised as an archetypal art image (which has lead to its appropriation and manipulation by other artists). What is even more well known than the image itself, is that it is a 'great' work of art, so its social cachet as art is even greater than its artistic value as a painting.

[2] I adopt Daniel Dennett's notation of capitalising the V to signify the sublime scale of the word used in this context; parts of this section draw their inspiration from a similar discussion on genotype–phenotype space in [2, Chap. 5].

Imagine if we were to iterate through every possible 500 × 400 pixel image, starting with all bits of each pixel set to 0 (the 'all black' image), changing bits one by one, until all bits for each pixel are 1 (the 'all white' image). In between the all 0s image and the all 1s image, would be a (partially) fascinating journey, because if we were able to do this, along the way we'd see every image that has ever been, or ever will be taken by anyone, anywhere! Every great (and not so great) work of visual art is in there, past, present and future, as are images of political assassinations, nude celebrities (even ones that have never posed nude), serial killers, animals, plants, landscapes, buildings, every possible angle and perspective of our planet at every possible scale and all the other planets, stars, galaxies in the universe, both real and imaginary. Pictures of next week's winning lottery ticket, and of you holding that winning ticket.

There are pictures of people you've never met or seen before (even in pictures, although there are many pictures of you with them, even looking at pictures of you looking at them). There are pictures of you and me together with our arms around each other like we've been best friends for years (even if we've never met), and pictures of you as a child sitting on my knee while I read what looks like a copy of this book to you.

Pictures of you at every moment of your life from conception to your death. It's not just you: there are pictures of every person who has ever existed at every stage of his or her life, from atomic close-ups to long shots. There are even some group portraits of all your ancestors (although admittedly at 500×400 pixels it is hard to make out a lot of detail). Then there are pictures of people that have never existed, along with pictures of people in situations that have not happened to them in reality. Then, there are all of these images (and many more) with every Photoshop filter ever invented (even the expensive third-party ones and even ones that haven't been invented yet), applied in every possible combination! And that's just a tiny fraction. *Every possible image*. Here is the image version of Borges' *Library of Babel*.[3]

Within this library of all possible images, along with all these interesting images, are many more that are not so interesting. So along with the *Mona Lisa*, for example, are all the *Mona Lisa* copies with just one pixel different (there are $3.35 \times 10^{12}$, or *three trillion*, of these). Then there are the ones with just two pixels different, and so on. In some versions, only parts of the image can be recognised as the *Mona Lisa*. Many others are just abstract patterns or shapes; some just look like noise or random bits of colour. Clearly, for each image that we know to be 'interesting' there are a lot of others that are almost as interesting, and as we get more and more distant from the interesting ones there soon comes a point where they are clearly not as good, eventually bearing little or no resemblance to the initial, interesting image.

Nonetheless, being able to generate every possible image sounds like a good idea, one that we could make a lot of money from (imagine selling all those

[3] The story of an imaginary library containing all possible 410-page books [3].

nude celebrity images to tabloid newspapers without having to actually go out and take them!). However, before you rush off to start writing the program to iterate through this set, it is important to understand how big the number of all possible 200,000 pixel images is. Even though each image is relatively small, with millions of them easily stored on a modern hard disk, the number, or *phase-space* of all possible images is very big. At 24 bits per pixel (the standard for colour images) its about $9.5 \times 10^{1444943}$. How long would it take to iterate completely through this phase-space? Lets be optimistic and imagine that every particle in the universe[4] is a supercomputer and can compute one billion images per second (we'll conveniently ignore the problem of how we'd actually look at them). Each particle[5] has been computing images since the universe began. How many images have been computed in total, i.e., how far have we progressed from the all black starting image through all possible $500 \times 400$ pixel images since the universe began? Fifty percent? Ten percent? One percent? The answer is approximately $2 \times 10^{-105}\%$. Yes, 105 0s to the right of the decimal place; in practical terms, basically none! Here we see the Vastness of combinatorial explosion, which occurs in all sorts of problems, not just images. It seems our financial security from tabloid photo sales has been put on hold.

What if we were to simplify the problem? Reduce the resolution (even low-resolution images of nude celebrities might fetch a good price). Reduce the bit depth — black-and-white might be good enough. Would that make it possible to iterate through in a practical time? Unless you're willing to accept very tiny bitmap images, the answer is, unfortunately, no. If you were prepared to spend your entire working life looking at images at a rate of one per second, you should just be able to look at all possible $5 \times 5$ pixel binary images. Before you attempt this, here's what the *Mona Lisa* looks like as a $5 \times 5$ pixel bitmap (magnified ten times):

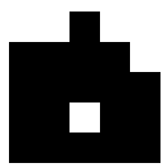

Pity the more ambitious fools who went one pixel higher in each dimension — they'd have only seen less than 1% of their set of possible $6 \times 6$ pixel binary images by the time they die.

It really is impossible to comprehend the size of this space of all possible images, even in relative terms, despite it being a finite set. Astronomical proportions, such as the size or age of the universe, don't come anywhere near to the measure of how big this space is. It is beyond the sublime, yet a computer

[4] For the purposes of this exercise, we assume there are $10^{80}$ particles in the universe — a reasonable approximation based on current estimates.

[5] I assume the number of particles in the universe is fixed over the lifetime of the universe; forgive me, this is only a thought experiment after all.

can generate any image from this set, so each image has the possibility of actually existing.

It is also interesting to observe that we could combine still images from the *everyimage* set in certain sequences to generate movies. So from the set of all possible $500 \times 400$ pixel images we could generate the set of all possible $500 \times 400$ pixel movies of some length. Since we have potentially every possible image, we also have every possible movie (including all the 'directors cuts', even if they were never made!). The catch is that to make a sequence we may need to duplicate some of the images (i.e., some of the frames might be the same).

This idea of duplication means that we could also build a $500 \times 400$ pixel image by tiling four $250 \times 200$ pixel 'quarter' images together. If our $250 \times 200$ pixel image set contains all possible images, this would include all the possible 'quarter' images from the set of possible $500 \times 400$ pixel images, the only condition being that in certain cases we will have to repeat some of the $250 \times 200$ images. For example, the 'all-black' $500 \times 400$ pixel image can be made by repeating the $250 \times 200$ pixel 'all-black' image four times. Why stop here? The set of all possible $125 \times 100$ pixel images could form the set of all $500 \times 400$ pixel images by using 16 of them at a time. If we follow this to its full regress, we end up with just two single-bit images: one containing a 0 and the other a 1. This is the binary universal image. It is capable of representing all possible images, and is easily searchable iteratively. The problem is that all possible images at this resolution collapse to either all black or all white, highlighting an important issue that will recur throughout this chapter, that of information and physicality. We need a certain amount of information (pixel resolution in this case) before we can physically start to recognise and distinguish images in some meaningful way. The resolution and recognition is dependent on the physicality of viewing — something that has evolved under constraints of efficiency, utility and fidelity [4]. Tiling or combining these 1-bit images together gives us all possible images at any resolution and bit depth; however, the difficulty is in knowing *which* bit to put *where*.

I hope I have now convinced you (if you actually needed any convincing) that the size of the search space for these types of problems is impractical for any kind of exhaustive search. The chances of randomly flipping bits with the hope of coming up with the *Mona Lisa* or even a nude celebrity are unimaginably *Small*. This is one reason why more sophisticated search methods, such as evolutionary computing methods, might be useful. But before we tackle that issue, there is one more important question to ask: Of the set of all possible images, what fraction would actually be 'interesting' images? That is, ones that we might actually want to spend some time looking at. Would this fraction be greater than or less than the fraction of 'junk' images (ones that we're not interested in looking at)?

Of course, interest is such an arbitrary thing at the micro level. I might be more interested in looking at pictures of my family rather than of a family I don't know. A medieval historian might be more interested in medieval castle pictures than modern architecture; a medical researcher might have a fancy for tumour images. However, these micro variations in interest don't matter statistically in the macro landscape of images humans create and have interest in. Additionally, what is interesting varies over time: you might find the first few fractal images you see interesting, but after seeing many fractal images your interest may wane — this behaviour characterised by the *Wundt curve* [5].

While it is difficult to pin down the exact number, it is clear that the fraction of interesting images from the *everyimage* set is extremely Small. If you need proof, try randomly generating $500 \times 400$ pixel images for a few hours and see how many interesting ones you find.

The problem of 'the possible and the actual' is well known in biology.[6] There are a large number of images that are actually interesting, but this set, even if a little fuzzy around the edges, is only a tiny fraction of all possible images. That's why random bits don't in general produce interesting images and why our brains are such good classifiers. The question is, can we automate the classification of images that are actually interesting from all possible images?

## Acknowledgements

I would like to thank the editors for their advice and feedback on earlier versions of this manuscript. I am also indebted to the following people for valuable discussions: Lloyd Allison, Peter Bentley, Jon Bird, Tim Blackwell, Paul Brown, Alan Dorin, Alice Eldridge, Mark d'Inverno, Janis Jefferies, Gordon Monro, Jane Prophet and William Latham. Parts of this manuscript were completed while I was a visiting researcher at Goldsmiths College, University of London. I would like to thank Prof. Mark d'Inverno for hosting my visit.

## References

1. Dreyfus, H.L., Dreyfus, S.E., Athanasiou, T. (1986). *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer.* Free Press. New York
2. Dennett, D.C. (1995). *Darwin's Dangerous Idea: Evolution and the Meanings of Life.* Simon & Schuster. New York
3. Borges, J.L. (1970). The library of babel. In Yates, D.A., Irby, J.E., eds.: *Labyrinths.* Penguin Books. Harmondsworth
4. Dennett, D.C. (1991). Real patterns. *Journal of Philosophy*, **88**: 27–51

[6] Many different DNA sequences are *logically* possible, fewer *physically* possible; a significantly smaller number still actually exist in biology.

5. Berlyne, D.E. (1971). *Aesthetics and Psychobiology.* Appleton-Century-Crofts. New York, N.Y.