Making art using evolutionary algorithms and artificial AI

T.D. Gedeon and J.Y. Shen
Department of Computer Science
Australian National University
Acton ACT 0200
Australia
tom@cs.anu.edu.au

Abstract

Evolutionary algorithms mimic the way genes mutate and recombine in environments which favours some individuals over others. Over time, the good components (genes in real, biological organisms) end up collected together to form better individuals while mutation (random minor changes) provides potential new good components. We need some way to evaluate each individual to determine whether it is worth keeping (or would survive in the wild), or not. But how do we do this with art, which is not readily analysed by computers?

This is where artificial AI comes in. AI is "artificial intelligence", which is really replacing (or replicating some interesting behaviour of) a person. Thus, "artificial AI" is replacing the computer by a person. So if an evolutionary algorithm encoded the components of an abstract computer generated picture, a human could identify 'nice' and 'not nice' images repeatedly to generate some art, which is tuned for their esthetic sense.

The art of Piet Mondrian is particularly suited for such experiments, firstly because of its apparently simple structure, and secondly as no simple mathematical formula has been able to be deduced for his work.

Introduction

In this paper we briefly introduce the reader to the work of the artist Piet Mondrian, discuss the computer generation and evaluation of art, propose our technical approach, describe our results and come to some conclusions.

Mondrian

> Pieter Cornelis (Piet) Mondriaan, after 1912 Mondrian, (pronounced: Pete Mon-dree-on, IPA: [pit 'mɔndɹiɔn]) (b. Amersfoort, Netherlands, March 7, 1872 — d. New York City, February 1, 1944) was a Dutch painter.
>
> He was an important contributor to the De Stijl art movement and group, which was founded by Theo van Doesburg. Despite being well-known, often-parodied and even trivialized, Mondriaan's paintings exhibit a complexity that belies their apparent simplicity. He is best known for his non-representational paintings that he called "compositions", consisting of rectangular forms of red, yellow, blue, white or black, separated by thick, black rectilinear lines. They are the result of a stylistic evolution that occurred over the course of nearly 30 years and continued beyond that point to the end of his life. (wikipedia, Piet Mondrian)

An example of one of Mondrian's abstract compositions is shown below:

Figure 1. *Composition of Red, Blue and Yellow.*

There have been multiple attempts to perform mathematical analyses (Hill, 1968, Reynolds, 1995) of the compositions by Piet Mondrian. None of them are successful on giving a convincing result revealing the "hidden math" within Mondrian's painting (wikipedia, de stijl).

Hill (1968) used "number math" for measuring the grid size, analyzing the ratio between grids, and so on. One example of the conclusions is that some of Mondrian's compositions are triple connected, in that you can not separate the graph into two without cutting at least three lines. This applies to about half his work in the period 1918 – 1938.

More success was achieved by Reynolds (1995), using "structural analysis" based on graph theory, which is correct for many of Mondrian's works, though certainly not all.

Art from artists or computers?

Mondrian's artistic role in esthetic choices in his compositions is still debated. Lee (2001) found that art students could not correctly identify genuine Mondrian compositions. Contrarily, McManus (et al, 1993) found that the majority of subjects could distinguish between original and modified Mondrian compositions. Wolach (2005) found that subjects could distinguish (preferred) Mondrian line spacings from divergent spacings. If subjects could select divergently spaced pictures that they preferred, then the preference for Mondrian spacings vanished.

Our view is that the differences are based on individual esthetics, which is our focus. That is, instead of identifying Mondrian compositions, we are interested in users constructing (with the aid of our tool *darwindrian*) Mondrian-like compositions which they find esthetically pleasing. We leave for later any investigation of any notion of general esthetic appeal of Mondrian's own compositions.

We also do not consider attempts to calculate esthetic worth of images, for example by Maiocchi (1991), and Garza and Lores (2005).

Generating Mondrian-like images

We construct Mondrian-like images using a number of parameters describing a possible image. These parameters are chosen by a random process initially. Subsequently, we use an evolutionary algorithm to improve the images for each user of our program.

Rules for Mondrian-like Graphs

In *de stijl*, only vertical lines and horizontal lines are allowed in the graph, and all lines terminate on other lines or the edge of the painting. The rectangle is also a basic element of Neo-Plasticism, but from a programming viewpoint, rectangles are a 'byproduct' of horizontal and vertical lines. Thus a Mondrian-like graph could be deemed a collection of horizontal and vertical lines.

We illustrate our algorithm below:

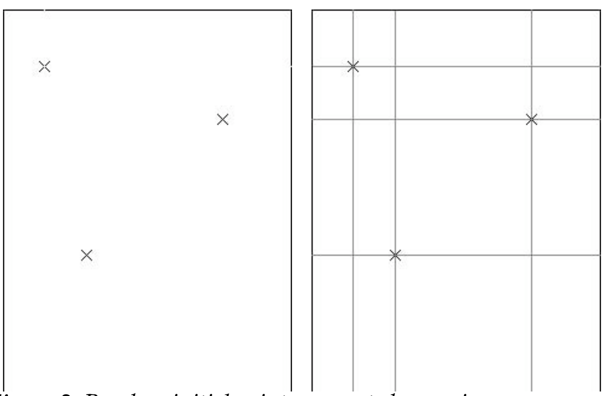

Figure 2. *Random initial points generated on a given canvas (left). Imaginary lines drawn crossing the initial points (right).*

Clearly there are many potential choices here. How many initial points do we generate? How close may they be to each other? How far apart? How close to the edges can they get? And so on. We can deduce many properties which often hold for Mondrian's own compositions. We have already mentioned his spacing of lines. Another example: Taylor (2003) analysed the positions of 170 lines featured in 22 paintings, and found that Mondrian was twice as likely to position a line close to the canvas edge as he was to position it near the canvas center. As we are interested in the development of individual esthetic choices, and we know from Wolach (2005) that the ability to chose can swamp the effect of Mondrian's own choices, we do not impose such conditions on our generation process.

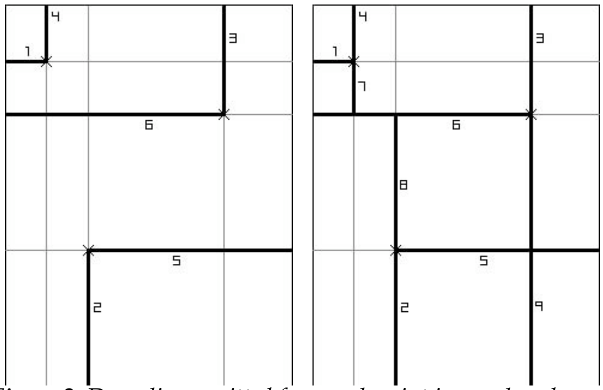

Figure 3. *Draw lines emitted from each point in numbered sequence (left). Skeleton complete (right).*

In the left part of Figure 3 (above), we show a partially completed skeleton. The emission of lines is probability based, so the same initial points can lead to different possible final images. Once all the lines are generated a remediation stage adds lines to eliminate remaining right angles, though this was not necessary in the example shown.

The final step is to randomly fill some rectangles with colour (red, yellow, or blue).

Evolutionary algorithms

Evolutionary algorithms are programs which solve 'problems' by simulating Darwinian selection among solutions. Solutions to a problem are represented in an abstract fashion, in a sequence of components of the solution. This sequence is called genes, and make up a chromosome by analogy with biology. An initial population of potential solutions (individuals) is generated by some random process. Each individual is evaluated, and the most fit are modified (recombined and possibly mutated) to create the next generation. Over time, individuals accumulate good components and are better solutions. In our work here, this would mean these individuals represent (can generate) 'nicer' Mondrian-like images, as they are the 'descendants' of a number of generations of better images.

In our program *darwindrian*, the chromosome consists of genes for the 3 parameters to generate our Mondrian-like images.

Parameter 1. Complexity

The complexity value is the most crucial value that affects the structure of the graph on which the image is constructed. It directly controls how many lines will exist in the final graph. The maximum lines in the graph (not including canvas edges) will not exceed 4 x complexity.

Parameter 2. Order vector

During the construction process, lines are drawn through the original points. This is a order-related process, for the latter lines drawn need to end on the lines drawn previously (to generate Mondrian-like images). Thus graphs can have the exactly the same original points but a different iteration order, to produce results which could be very different (although they may share a similar layout). Suppose there are three original points a, b, c, an order vector is an arrangement of a, b, c with each element appearing 3 times (the complexity) and the total length of the vector being 9. If the original points in Figures 2 & 3 were labelled a, b, c counter-clockwise from the top, then the order vector for the graph shown in Figure 3b is [abcabcabc].

Parameter 3. Structure vector

When a line is drawn from an original point, there are limited choices for its direction: North, South, East, West, since Mondrian-like graphs allow only horizontal or vertical lines. We assign a structure vector for each original point, which contains probability values for each particular direction to be selected. For example: [N, S, E, W] = [15%, 15%, 50%, 20%]. Since this is a random process, directions with less probability may still be selected. Taking this approach the program avoids being deterministic, thus showing some minor creativity (or at least an illusion of it). After a direction is taken, its probability value will be set to 0 to avoid being selected again, and other directions will enlarge in accordance to their original ratios so that the sum of probabilities in the structure vector remains 1. For example, after the E direction is taken, the example structure vector will become: [N, S, E, W] = [30%, 30%, 0%, 40%]

Fitness function

Evolutionary algorithms require a fitness function to evaluate the successfulness of each individual. In the first generation they are created with random properties, but subsequently they are constructed from the components of successful individuals. We use human assistance for this evaluation, and discuss this in the next section.

The best known related work is by Sims (1993), in which users "stand on sensors in front of the most aesthetically pleasing images to select which ones will survive and reproduce to make the next generation".

We note that in general "evolutionary computing approaches need to rely on fast execution times and the use of human interaction in the evaluation loop is simply too slow"(Punch, 1999). We compensate by maintaining a cumulative fitness model which is updated every time a user selects an individual (indicates s/he likes one of the Mondrian-like generated images). It is this cumulative model which is used to evaluate all members of a generation, hence reducing the impact on users.

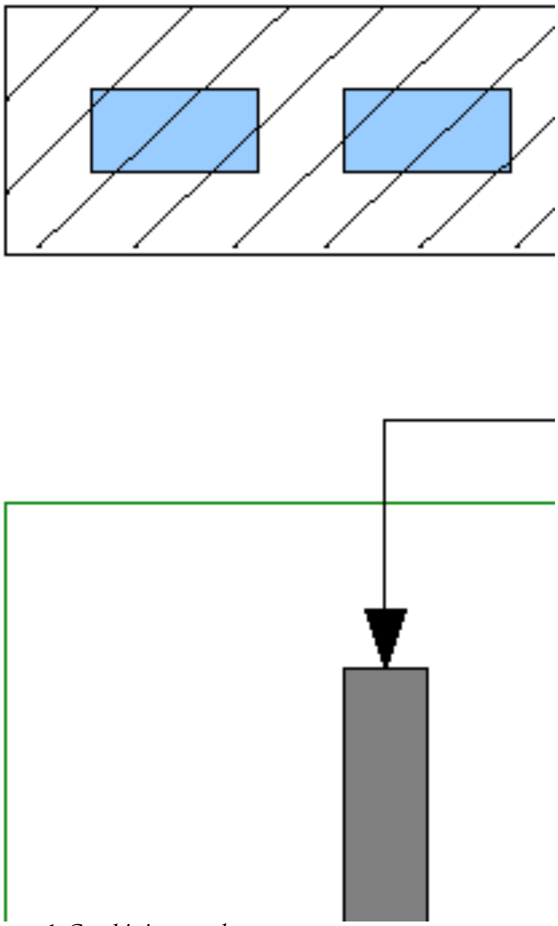

Figure 4. *Combining good genes*

Artificial AI

Amazon's Mechanical Turk was launched in late 2005. The site allows "software developers and businesses the power to use human intelligence as a core component of their applications and business" (Barr and Cabrera, 2006). The name refers to the famous 18th century (fake) chess automaton which had a human dwarf hidden inside.

Critics complain that the network is no more intelligent than its smartest members, and that it is a virtual sweat shop (Pontin, 2007). While the tasks being solved are well below the average capabilities of the members the first criticism has little impact, and people will contribute while there are substantial differences between earning power between poorer and richer countries.

A worse problem is the use of such artificial AI to defeat computerised challenges to determine whether the user is human or not. Thus a "completely automated public Turing test to tell humans and computers apart" (Wikipedia, CAPTCHA) can now be defeated by humans for computer software, albeit not in real time.

Arguably, Google's "google answers" which ran from 2002 to 2006 is a precursor of this concept.

Crowdsourcing (Howe, 2006) is perhaps a more accurate term for artificial AI, and appears more common in the artistic community.

A well known example is the Sheep Market (Koblin, 2006), though arguably this is not an artwork which uses crowdsourcing, rather it may be a compilation using the internet. See Martineau (2007) for a list of collaborative visual arts.

In our work we require the user to participate in providing the fitness function ('niceness' evaluation) for our evolutionary algorithm. We can however make the distinction that in this case the user is performing acts for their own direct benefit (as opposed to indirect benefit in the form of minor financial rewards). Hence it is possible that we should instead consider our case to be one where we are setting the preferences of a program, which does not have simple selectable choices but we must do so by example. Thus, we could characterise our technique as "implicit customisation".

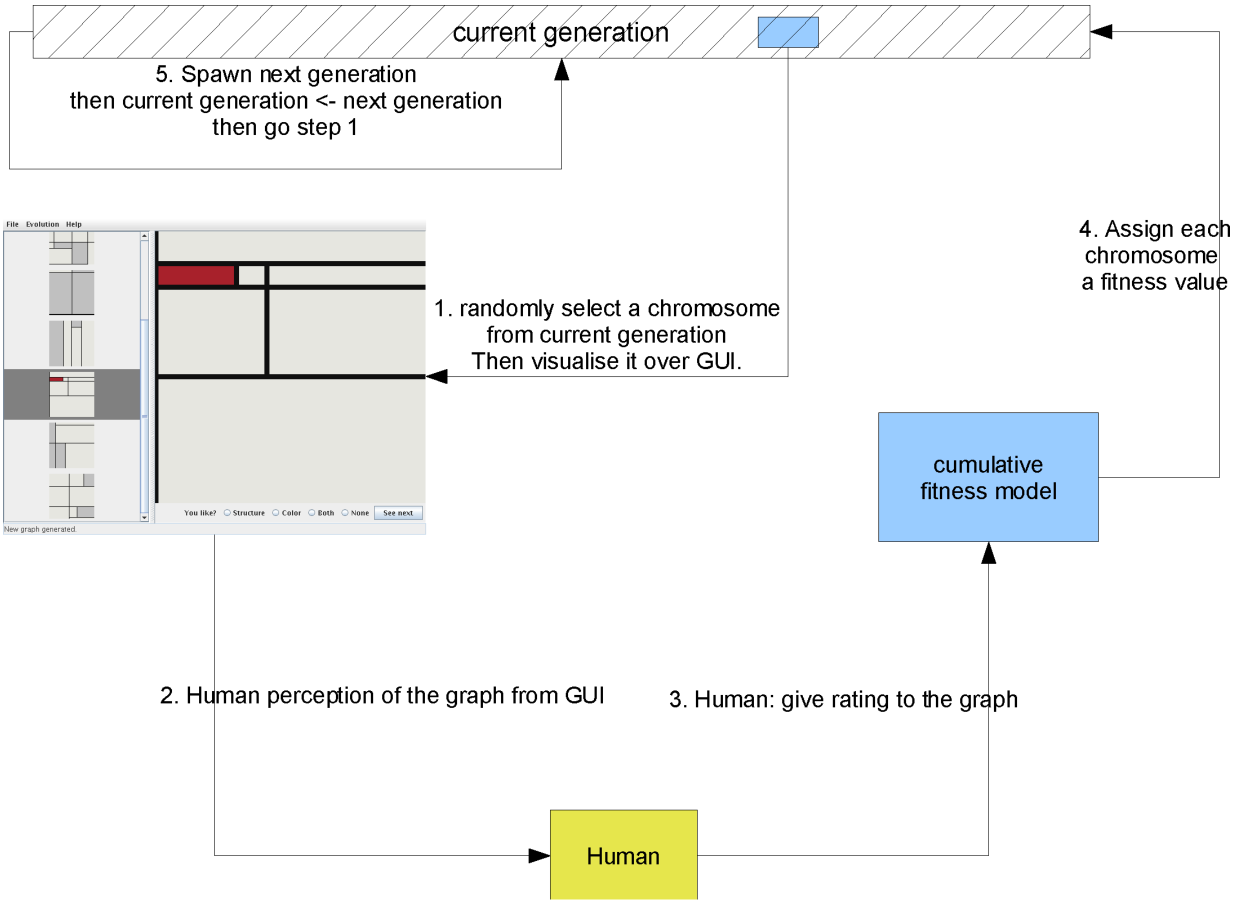

Figure 5. *Graphical User Interface (GUI) and user process diagram*

Results

We have generated Mondrian-like images for 6 subjects (Z, P, M, B, J, and T) using our *darwindrian* prototype.

A representative sample of 4 of the final 20 images for each subject is shown in Figures 5 to 10. We make some preliminary comments on the differences.

The subject (Z) was the most consistent in the choices made, as the key decision criteria can readily be seen even from the sample of 4 images.

Clearly, for Z's esthetic both red and blue must be present, and be (mostly) touching along a rectangle boundary.

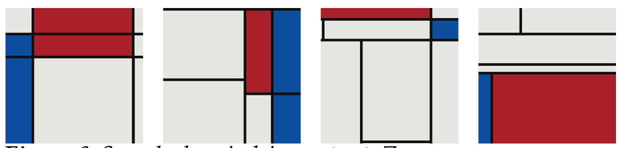

Figure 6. *Sample darwindrian output: Z*

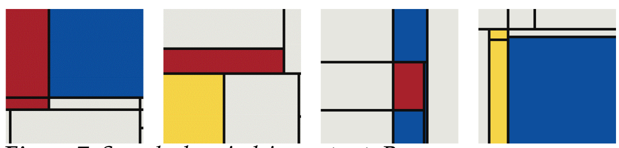

Figure 7. *Sample darwindrian output: P*

Subject P seems to prefer 2 colours, particularly including blue, which usually touch, and extend across at least one of the horizontal or vertical dimensions.

Figure 8. *Sample darwindrian output: M*

Subject M appears to favour the colour yellow, particularly in combination with red. The sizes of blocks of colour also are more consistent than for other subjects.

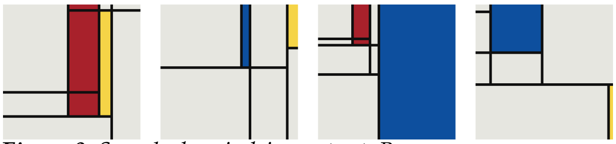

Figure 9. *Sample darwindrian output: B*

We could not easily deduce what subject B prefers.

Figure 10. *Sample darwindrian output: J*

Subject J clearly prefers small blocks of colours very close to the edges.

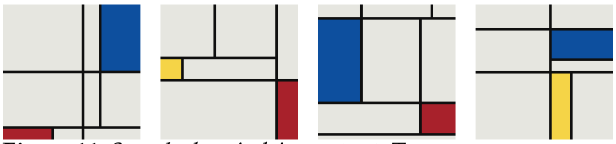

Figure 11. *Sample darwindrian output: T*

Subject T seems to prefer 'medium sized' blocks of colour separated by some space.

Subjects M, J and T were familiar with Mondrian's paintings prior to this study.

Conclusion

In the previous section we showed the results from our *darwindrian* prototype for 6 subjects. We found that the images for each subject were quite different, and it was generally possible to hypothesise a decision rule for the choices made by each subject.

We can therefore conclude that it is possible to generate Mondrian-like images, which correspond or are tuned to individual esthetic appreciation.

In our future work we intend to analyse the differences between images liked by different users and attempt to speed up and direct the evolution of images more pleasing to new users. Success in creating such an automatic fitness function would be a first step to an encoding of an esthetic function on these images.

We also intend to evaluate differences between conscious choice versus attention in the evolution of our images, by the use of an eye gaze detector and EEG recorder.

References

Barr, J. and Cabrera, J. F. 2006 "AI gets a brain: New technology allows software to tap real human intelligence," *ACM Queue*, 4(4), May 2006.

Garza, A. G. and Lores, A. Z. 2005. "Evolutionary Art Revisited: Making the Process Fully Automated," *Proceedings WSEAS International Conference on Soft Computing (SOSM)*, pp. 73-79, Cancún, México.

Hill, A. 1968. "Art and Mathesis: Mondrian's Structures," *Leonardo*, I, pp. 233-234.

Howe. J. 2006. "The Rise of Crowdsourcing," *Wired*, June 2006.

Koblin, A. 2006. "The Sheep Market," www.thesheepmarket.com

Punch, W. 1999. "Review of Genetic Programming 1998," *IEEE Transactions on Evolutionary Computing*, 3(2), pp. 159-161.

Lee, A. 2001. "Piet Mondrian: Buridan's Ass and the Aesthetic Ideal," *Proceedings of the International Conference on "The Art of Seeing and Seeing of Art,"* Australian National University, (Canberra).

Locher, P. J. 2006. "The usefulness of eye movement recordings to subject an aesthetic episode with visual art to empirical scrutiny," *Psychology Science*, 48(2), pp. 106 – 114.

Maiocchi, R. 1991. "Can you make a computer understand and produce art?" *AI & Society*, 5: pp. 183-201.

Martineau, J. 2007. "Report on the online art community deviant art," http://tinyurl.com/22jt9v

McManus, I. C., Cheema, B. and Stoker, J. 1993. "The aesthetics of composition: A study of Mondrian," *Empirical Studies of the Arts*, 11(2), 83-94.

Pontin, J. 2007 "Artificial Intelligence, With help from the humans," *New York Times*, 25 March 2007, http://tinyurl.com/2pl49v.

Reynolds, D. 1995. *Symbolist Aesthetics And Early Abstract Art*, Cambridge UP, page 260.

Sims, K. 1993. http://www.genarts.com/karl/genetic-images.html

Taylor, R. P. 2003. "Fractal expressionism-where art meets science," *Art and Complexity*, Elsevier Press (Amsterdam).

Wikipedia, entries on “de stijl,” “Piet Mondrian,” “CAPTCHA,” http://www.wikipedia.com

Wolach, A. H. 2005. “Line spacing in Mondrian paintings and computer-generated modifications,” *Journal of General Psychology*, July, 2005.