Evolutionary and Adaptive Systems MSc.

Creative A.I. Term Paper

An Investigation into Aesthetic Preference using Generative Art.

Mike Blow
Candidate Number 80761

1.0 INTRODUCTION

The existence of a universal aesthetic preference has long been a topic of debate. Several years ago, thinking about generative art, I realised that it would be possible to evolve pictures according to the viewer's personal preference. I also realised that the resulting pictures, if taken from a wide variety of people, might indicate some kind of universal aesthetic preference. This paper describes a 'first-draft' implementation of that idea. Prior to that I present an overview of the study of aesthetics and of computer-generated art.

2.0 GENERAL BACKGROUND

2.1 Aesthetics

Aesthetics can be defined as the philosophy of beauty. In 'The Critique of Judgement', Kant argues that aesthetic judgements must be disinterested, universal, necessary and final. In other words an object should give us pleasure because it is beautiful and not the other way around, it should be accepted as such by the majority and it should affect us as if it had a pupose even if one does not actually exist. The two extreme views of aesthetic study are rationalism, which analyses various attributes of an object to study its beauty, and romanticism, which views a work holistically and focuses on the instinctual and emotional response to it. Often a combination of both approaches can provide us with the best analysis as each approach complements the deficiencies in the other.
An object can be called beautiful for many reasons. Definable physical qualities such as composition, colour, harmony, symmetry, and fractal value are the most obvious, but there are also emotive judgements such as powerful, restful or disturbing. Finally there are criteria that may be tangential to the physical appearance of the work: fashion, politics, world events, functionality and social trends. Judging beauty is intuitively easy; defining it is very difficult.

2.2 Some Aesthetic Attributes

The composition of a painting, position of the subject matter, proportion of the elements in the frame and indeed the shape of the canvas have an important aesthetic effect. Rules of composition like the 'golden ratio' (about 1.618) are used as guidelines to what intuitively looks pleasing to the eye. However the skillful artist can create interest using assymetrical compositions, such as Degas' paintings of ballet school in the 1870s.
The tension between order and disorder is also a fundamental concept. Balanced composition and harmonious use of colour are generally thought to add to the beauty of a painting, but too much regularity can cause loss of interest. Novelty and complexity are desirable to gain, and hold, the attention of the viewer. It is worth noting that complexity can appear on many levels. For instance, the initially simple shapes and limited colours of Mark Rothko's work exhibit, upon closer examination, complex subtleties of tone and texture deriving from many layers of brush strokes.
Fractals, the phenomenon whereby a similar form repeats at different scales, have been studied since the turn of the 20th century, but were not formally defined until the 1970's [3]. A study of the drip paintings of Jackson Pollock has shown that they increased in fractal dimension throughout his career so consistently that it would be possible to date them from this value alone [10]. They contain a type of aesthetic not measurable without detailed examination, but presumably the celebrated pattern recognition capabilities

of the brain cause this to be found attractive on a more subconcious level.
Symmetry is so naturally prevalent it is easy to take for granted, but it plays a major role in composition. Designs and objects that are symmetrical are often more pleasing to the eye than assymetrical ones. Symmetry implies balance, efficiency, order and a link to the familiar natural world.

2.4 The Aesthetics of Nature
Most people appreciate the colours of a beautiful sunset, the grandeur of mountains or the grace of dolphins. It is fair to argue that the roots of our aesthetic appreciation derive, at least in part, from nature. Natural objects show a large fractal dimension and high degree of symmetry. A coastline, cloud or fern are all objects that repeat at many different scales. Twofold symmetry is found in almost all creatures, but six-fold symmetry occurs in snowflakes and even four and five-fold in flowers and seeds. It has been shown that humans have a distinct preference for abstract images that approach the fractal dimension of nature [8].

2.5 The Universal Aesthetic
Beauty may be in the eye of the beholder but do all eyes share some aesthetic preferences? It is of course impossible to decide whether some object definitely posseses beauty as appreciation of that beauty is a subjective response; however it is not impossible to ascertain that to the majority of people it is aesthetically pleasing. A good example given by Sudweeks and Simoff in [9] is the Taj Mahal in India. The general consensus is that this is a beautiful building, and the number of people of all nationalities and cultures who come to visit it strengthens the idea. Nevertheless, the universal aesthetic is not universally acknowledged. Some argue that aesthetic judgements are entirely a product of culture and prior experience, and thus that only something familiar but novel will be judged as interesting [4]. It is suggested that something too well known will be discarded as boring, but that something too new will confuse. Aesthetic development must therefore progress in small steps in order to be accepted.

2.6 Modelling Aesthetics
Computers do not generally lend themselves to modelling philosophical concepts. However if the rationalist view is adopted a computer can be put to good use measuring attributes of a work. Computer art is also often easily created and manipulated and so lends itself to ‘large scale’ comparison studies of aesthetic concepts.

2.7 Computer Generated Art
Computers have been used to create art for many years. Currently computer art systems can be roughly divided into three types. There are those that mimic the work and methods of a human artist, of which Aaron by Harold Cohen [1] is a celebrated example. Aaron uses a representational world model approach to draw ever-changing views of the interior of a room with people and plants. It has been developed over many years and initially imitated a child’s drawing process. Over time it became more complex and expanded its repertoire to include plants, 80’s fashions and humans with strange beards.
Some artists explore a conceptual space using algorithms. The algorithmists produce intricate patterns deriving from mathematical formulae, an approach only really practical with the advent of computers. The art does not really mimic the human painting process (although it is in a sense an extension of the geometric art of escher, for instance).
Lastly there are programs that help a human to explore new artistic territory through the use of evolutionary processes. Dawkins’ biomorphs [2] allowed the user to evolve natural-looking forms. Each shape was defined by 9 values held in a genome, and the current shape plus 8 mutations of it were presented to the user. By cliking on his favourite, the user could direct evolution to produce a large variety of shapes. Sims [5,6] applied the idea to computer graphics and evolved abstract images, and Latham and Todd [12] combined the two to produce organic forms rendered in three dimensions.

Thomas [11] also used directed evolution to evolve three-dimensional forms. Typically these programs work by defining the parameters of the form or picture in a genome of real or binary values. A series of images is created from random variations of the genome and the results displayed on screen. The artist chooses one of the pictures according to his preference and is then presented with more pictures, based upon mutations to the genome of the chosen picture; he chooses one of those and is presented with more options and so on. In this way the artist can explore a large conceptual space very quickly and react intuitive to what he finds aesthetically pleasing.

It should be noted that in none of these approaches does the computer learn any aesthetic values. In each case it is essentially a tool, an extension of the paintbrush, to allow new and hitherto perhaps impossible images to be created, but relying on the aesthetic preference of the user for value judgements. In the case of Aaron it is primarily a model of the human creative painting process but the omission of learning means it will never progress or improve. An enormous part of the human creative process relies on iteratrive learning and value judgements and in this resepct Aaron falls far short of being a convincing model.
Obviously attempting to model human aesthetic value judgements is very difficult, although Saunders has modelled curious agents which react to novelty, an essential part of the exploratory artistic process. It has been argues that visual aesthetic appeal depends in part on the surprise generated by novelty, and as such this could be seen as the first step towards an aesthetically-aware agent.

3.0 AN EXPERIMENT IN EVOLUTIONARY ART

3.1 Overview

My program CAT (Common Aesthetic Test) is a generative art system using a simple set of parameters to produce a variety of abstract pictures. I decided to try and remove as many elements that might cause 'pseudo-aesthetic' judgements as possible – partisan judgements or decisions based on current fashions or events – and reduce the images to simple shape collages. Having said this the pictures produced are stylistically reminiscent of 60's and 70's abstract art. I hoped that making the pictures abstract would allow me to study people's preferences to fundamental attributes like symmetry and complexity while sidestepping personal biases caused by subject matter. Nevertheless there is enough variation in these simple constraints to produce a wide range of interesting forms.

3.2 Picture Parameters

A picture is formed of an 8x8 grid of shapes. A single type of shape, the 'rounded rectangle', is used. This shape is useful in that it can, with the modification of only 1 roundness parameter, create a range of shapes from a perfect square (or rectangle) to a perfect circle (or oval) and anything in between. All the parameters affecting these shapes are defined in the pictures genome. Some of the pictures created by CAT are shown in Appendix A, showing the range of images possible from the limited parameter set.

Base Width
Description: The initial width of the shapes before any variations are applied.
Range: 0-300
Mutation Range: +/- 15

Width Variation
Description: The maximum variation (positive or negative) applied to the basic width.
Range: 0-300
Mutation Range: +/- 15

Base Height
Description: The initial height of the shapes before any variations are applied.
Range: 0-300
Mutation Range: +/- 15

Height Variation
Description: The maximum variation (positive or negative) applied to the basic height.
Range: 0-300

Mutation Range: +/- 15

Roundness
Description: The roundness of the corners of the rectangles. These can vary from fully rounded to perfect right angles.
Range: 0-400
Mutation Range: +/- 20

Number of Colours
Description: Range of colours from which the program chooses randomly. There is a maximum of 8 predefined colours, but variations in transparency and overlaid shapes can produce many more tones.
Range: 1-8
Mutation Range: +/- 2

Transparency
Description: Transparency of the shapes.
Range: 0-10
Mutation Range: +/- 2

Fractal Value
Description: Fractal quality of the shapes. this was a simple 'on or off' parameter, with fractal shapes at different scales being created by adding the same variation to both height and width.
Range: 0-10 (0-5 non-fractal, 5-10 fractal)
Mutation Range: +/- 2

Symmetry
Description: This parameter has four settings: no symmetry, repeated pattren (the 4x4 grid in the top left corner repeated in each of the other quadrants without mirroring), two-fold symmetry around the vertical axis, and four fold symmetry. Note that the randomness of colour choice disguises any symmetry to some degree, unless only one colour is being used.
Range: 0-40 (0-10 no symmetry; 10-20 repeat; 20-30 twofold; 30-40 fourfold)
Mutation Range: +/- 4

3.3 Genotype to Phenotype mapping

An interesting problem in generative art is how to make successive pictures look neither too similar nor too different to their forebears. Pictures based on algorithms can look very different given only small changes in values, meaning images based on very similar genotypes can produce confusingly different results (or phenotypes). It is difficult to guide evolution if the mutations cause large changes in the phenotypes. The obvious solution is to define every aspect of the work to give a direct genotype to phenotype mapping but this has one major drawback, in that the genome for a large or complex picture could be enormous. Also artists tend to embrace serendipity and a certain sense of 'exploration' and for me, certainly, predefining every aspect of the work somehow seems to be ignoring a useful force in the artistic process.
CAT approaches this problem by defining in the genotype ranges from which values are randomly picked instead of specific values. Thus a picture where the basic width of a shape is 50 and the width variation range is 20 will always have widths randomly chosen between 30 and 70; and will always look similar, but not identical. It should be noted that this means 2 pictures made with the same genome will never (for all practical purposes) look identical but should appear to be from the same 'family'. Mutations upon the genome might then affect the width variation to produce a different aesthetic effect. Soddu has used a similar system in his Argenia evolutionary design program [7] where the random elements of the designs are seeded from the system clock meaning there will never be two identical products.

3.4 A Word about Genetic Algorithms

Normally genetic algorithms explore the search space of possibilities using both crossover (where the childrens' genes are a mixture from two or more parents) and mutation (where certain values are changed slightly from one generation to the next). As there is only one parent CAT uses just mutation. Indeed in this case crossover would be detrimental, causing 'child' pictures to to have too much variation for easy assessment. The other major difference in application is that GAs usually have some automatic fitmess evaluation and can often run many generations in a short time, whereas in this case the user's aesthetic judgement is the fitness function which limits the system to real time operation.

3.5 Method

Initially four pictures are created with random parameters (*fig. 1a*). The user selects their preferred picture and the next generation of four images are displayed based on mutations to the chosen genome. The genome for the top left image is unchanged, and the other three have mutations to the picture parameters (*fig. 1b*). The process is then repeated to evolve the desired picture.

In the test the subjects were told to choose the picture they found the most aesthtically pleasing in each generation. In the fifth generation, the genome of the chosen picture was saved to disc and four more random 'starting pictures' were displayed. At any time the subject could click on a reset button to start over with four random pictures. No time limits were imposed. In general the pictures created by the system seemed to interest people and some spent quite a long time exploring the images created by the system.

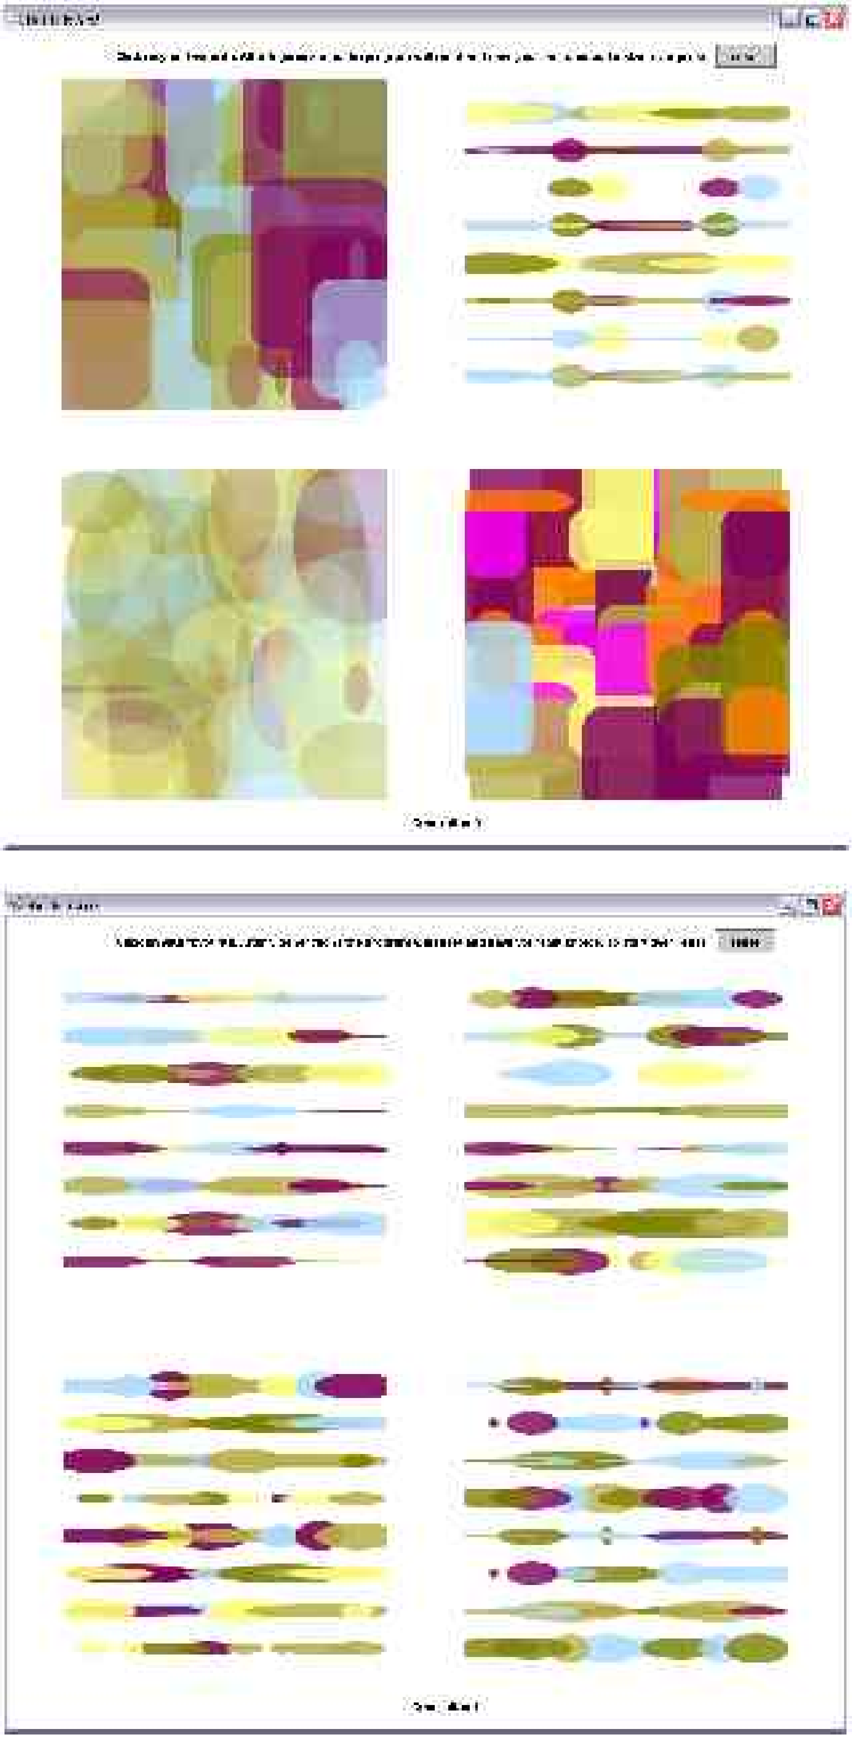

figs. 1a and 1b. Generations 0 and 1 of a typical run. The top right picture in generation 0 was chosen and mutated to give the four options in generation 1.

3.6 Results
127 results were collected by testing 22 subjects. Seven nationalities were represented. The first four results were fairly evenly spread with no real obvious preferences.

Base Width in pixels.
Average base width = 104

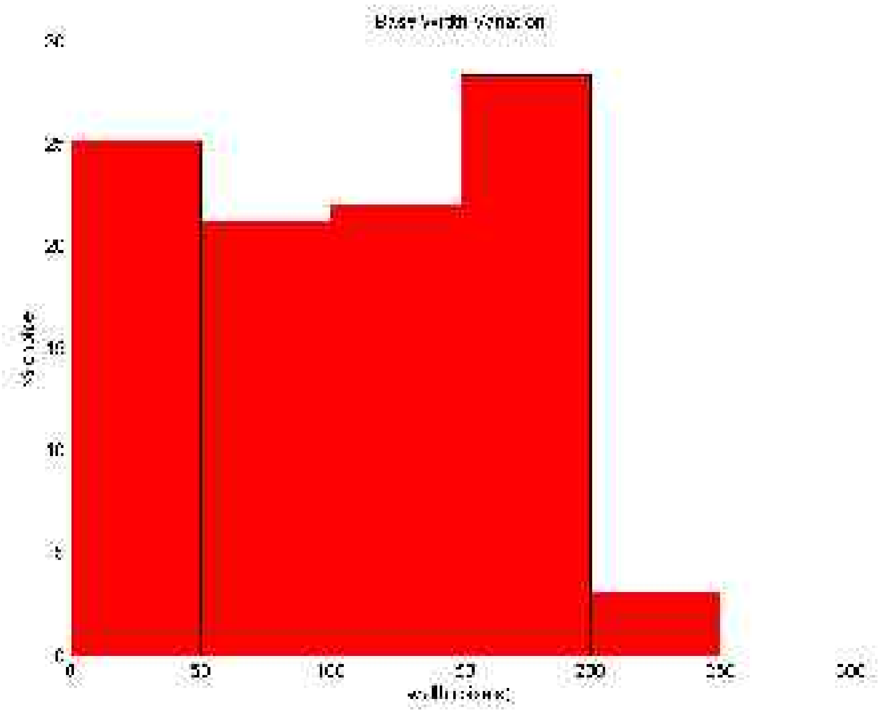

Base Height in pixels
Average base height = 94

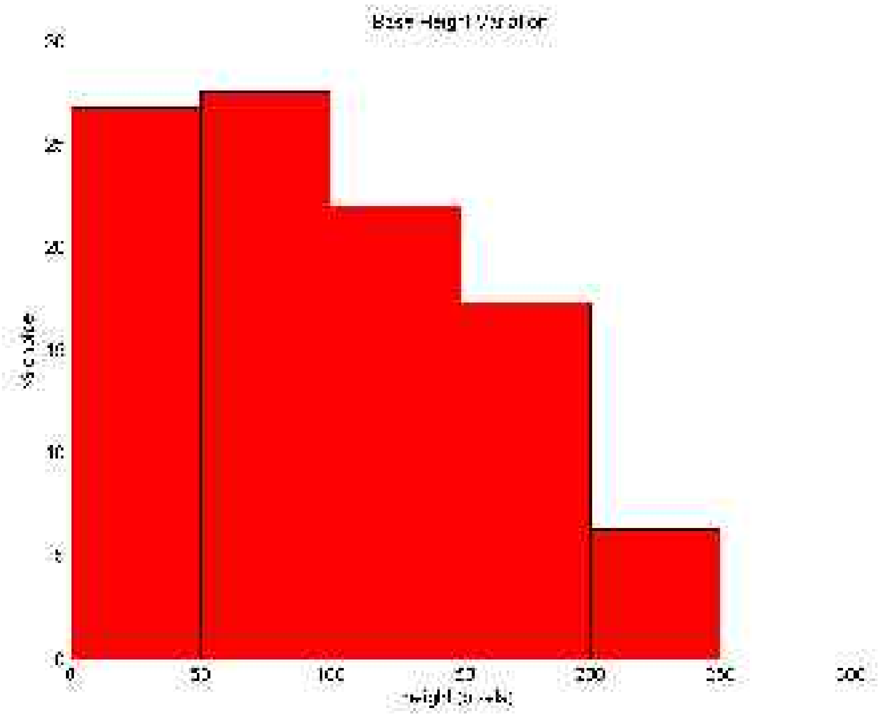

Width Variation in pixels
Average width variation= 94

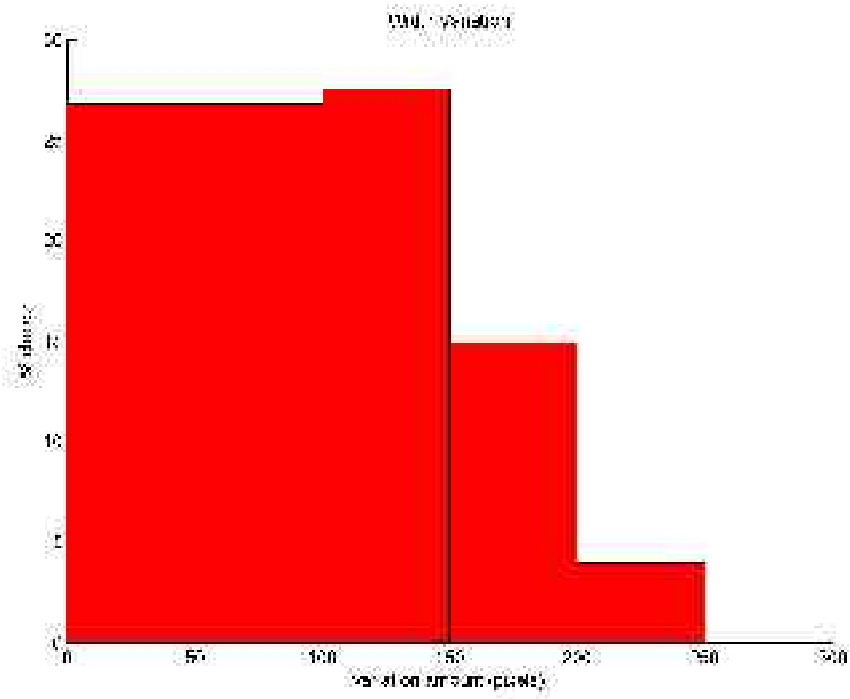

Height Variation in pixels
Average height variation= 98

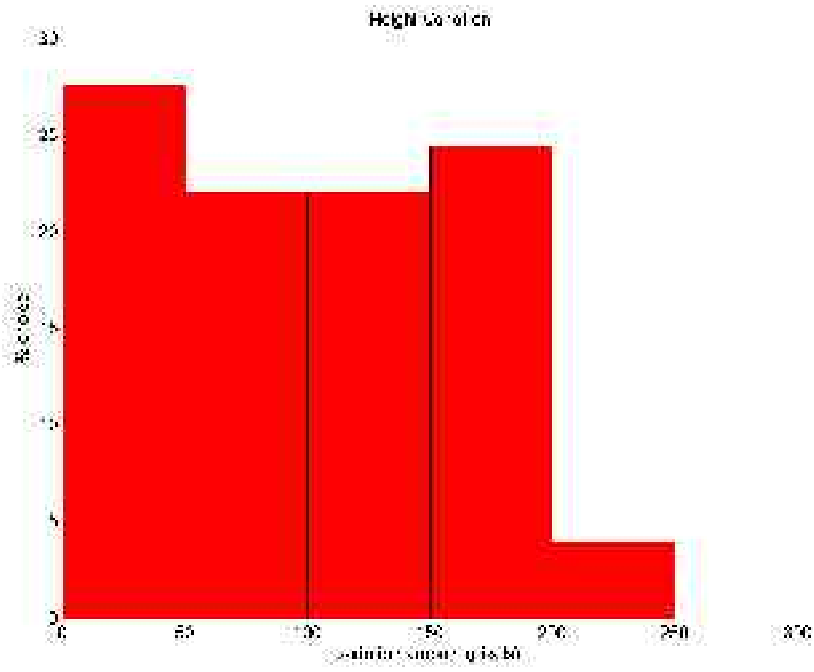

Roundness Value
Roundness is an absolute value so the amount of roundness seen depends on the size of the rectangle it is applied to. A small amount of roundness applied to a small shape will make it practically a circle, while a larger shape will remain essentially a rectangle or square. A small preference for low roundness was shown in these tests.

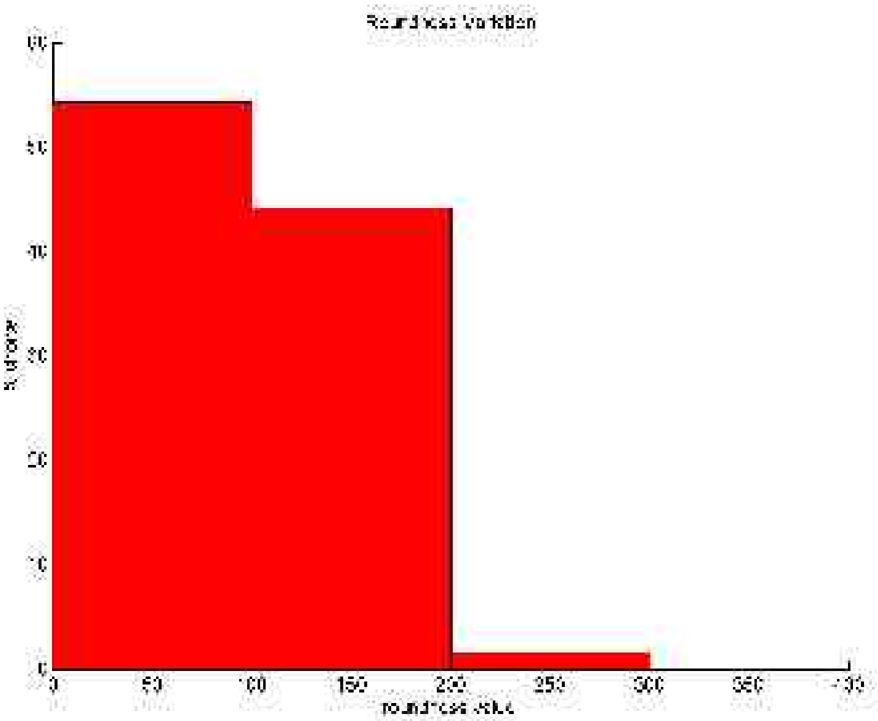

Number of Colours
General preference for 5 or more colours but interesting preference for 1 colour.

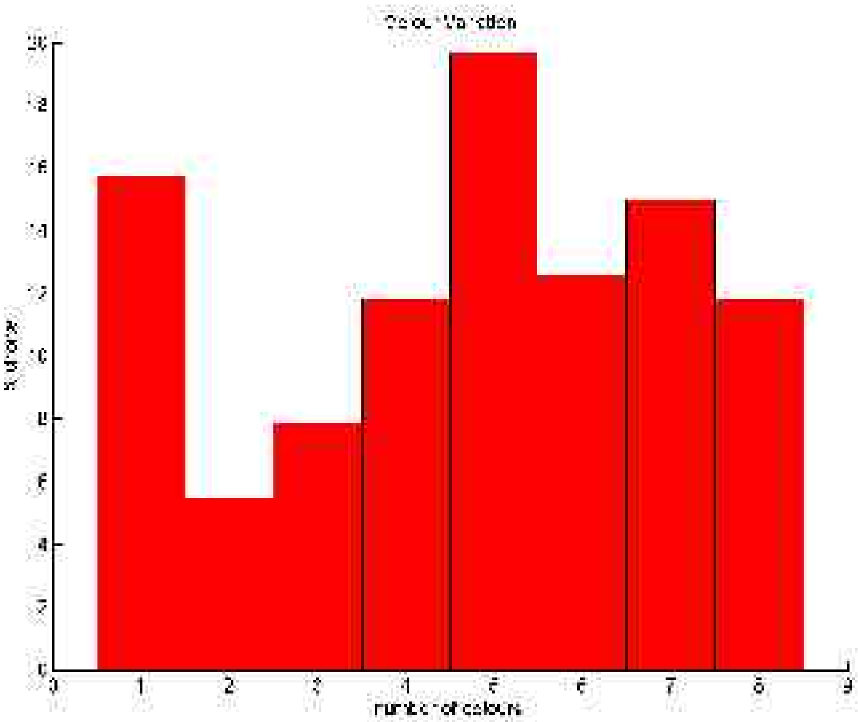

Transparency
Clear preference for low transparency.

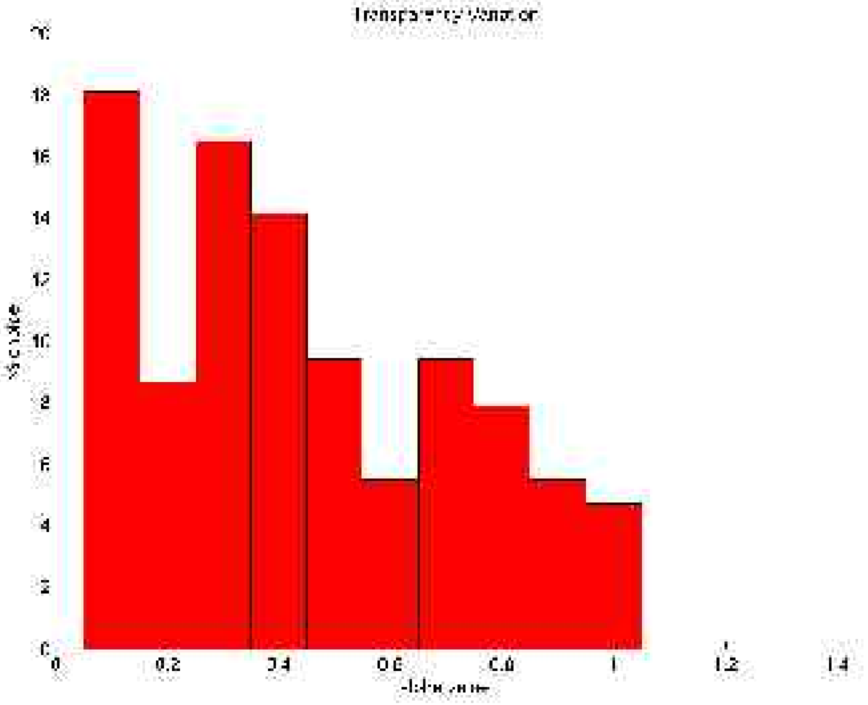

Fractal Value
No obvious preference.

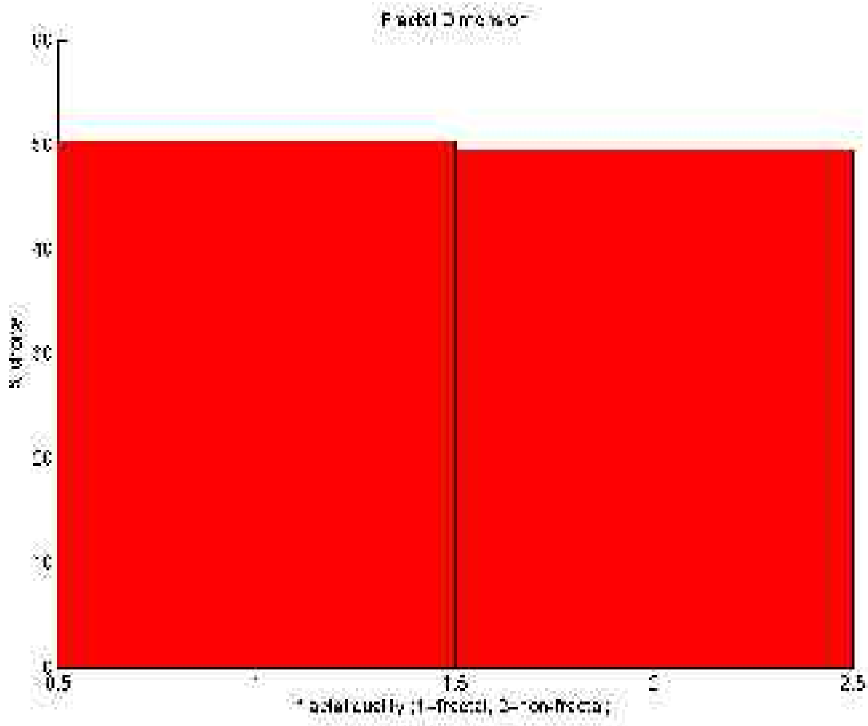

Symmetry
Twofold symmetry is preferred at about 30%.

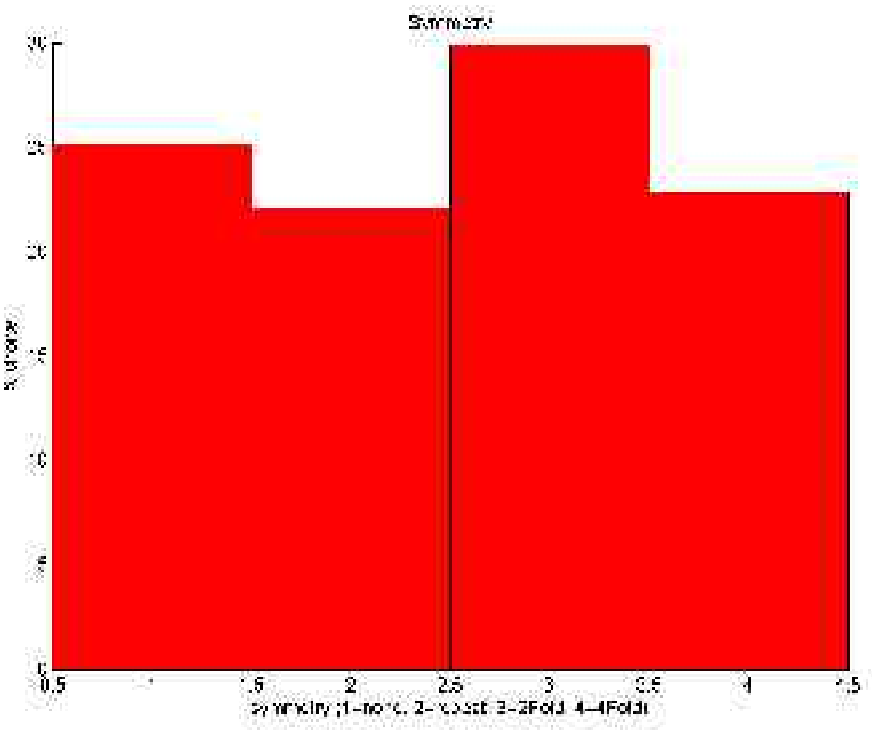

4.0 Discussion
Several results would seem to suggest that there is a common aesthetic preference amongst this test set. Twofold symmetry is preferred which is interesting as, from the options presented, it is the one most prevalent in nature. In addition it does not produce patterns as symmetrically complete as fourfold, so the choice would seem to be based on something other than aesthetic balance alone. There is a distinct bias towards high transparency. I would suggest this is because it adds two very important qualities to the picture; complexity and harmony. Overlapping semi-transparent shapes produce a mixture of familiar shapes (the shapes themselves) and random shapes (the overlapping areas) which I believe might form an attractive compromise between order and chaos for the viewer. I suspect this is why there is a marked and uncharacteristic preference for a single colour as well, in that it allows the viewer to more accurately focus on the form of the picture without distractions caused by other colours. Furthermore transparency gives the picture more depth. The nature of the drawing process means that some of the shapes are drawn in front of others and transparency gives the viewer a chance to appreciate the 'levels' of the drawing. Harmony is enhanced because the colours, being semi transparent pastel versions, complement each other better and combine to form more pleasing shades, in a 'stained glass window' kind of effect.
The other results are not particularly conclusive, being fairly evenly distributed. However I believe in part this is because the test subjects were often exploring the artistic possibilities of the system, and that sometimes the novelty value of the pictures overruled other aesthetic judgements. Perhaps some parameters were too hard to discern. I suspect the fractal value was overlooked because the user's attention was drawn to other aspects of the pictures. In terms of defining a picture the number of parameters used is quite small, but for a test situation it is relatively high. Simpler pictures would allow more accurate parameter measurement but it must be said

that pictures that were too simple might lack the depth to be aesthetically pleasing. In conclusion it is apparent that a certain level of complexity and novelty are a preferred choice in this group of subjects. This would seem to bear out the idea that an interesting picture needs to have a tension between order and chaos. It must be both novel enough to catch our attention and complex enough to hold it.

5.0 Further Work

CAT could be expanded in many ways. The current failings of the system are its localised nature, the fact that some parameters might 'mask' others and the inability to accurately reproduce a saved picture from the genome due to the 'random range' effect. To address the first issue a web version could be created. A saved image could be recreated exactly by saving the random seed as part of the genome. To enable more specific evaluation of aesthetic judgements, one could create a set of pictures with this system and analyse each for the various attributes, and then display these 2 at a time and ask the subject to choose between them. This would allow more accurate control and mean the pictures were constant across all subjects. The fractal value of the pictures could be calculated using proven methods [10] rather than having a simple fractal/non fractal split. It would also be interesting to test groups of people (perhaps various ethnic groups, men and women, designers and non-designers) to get some specific results.
Regarding the genetic element of the pictures, an idea occurred to have a flat screen on the wall of a gallery showing a generation of pictures one at a time with an 'interest feedback system' consisting of a camera to measure the amount of time people were looking at each picture. New generations of images could be bred from the most popular picture(s).

6.0 Conclusion

Intuitively it seems that people are not so diverse as to find beauty in completely different things, and that there must be some kind of universal aesthetic. However to prove this would take a carefully designed test administered to a wide cross section of the world's population. The CAT test described here has a very limited range, but has produced interesting indicators that there are some common preferences among a small group of subjects. As a first draft implementation it has also raised several issues that could be improved or investigated in the future.

REFERENCES

[1] Cohen, H. *The further exploits of Aaron, painter.* Stanford Electronic Humanities Review 4(2), 1995.
[2] Dawkins, R. (1987). *The Blind Watchmaker.* W.W. Norton and Company. New York, London.
[3] Mandelbrot, B. (1987). *The Fractal Geometry of Nature*, W.H. Freeman and Company.
[4] Saunders, R., Gero, J.S. (2001) *Designing for interest and novelty*, Proceedings of Ninth International Conference on Computer Aided Architectural Design Futures, pp. 725-738.
[5] Sims, K. (1991). *Artificial Evolution for Computer Graphics,* Computer Graphics (Siggraph '91 proceedings), pp.319-328.
[6] Sims, K. (1994). *Evolving 3D Morphology and Behavior by Competition*, Artificial Life IV Proceedings, ed.by Brooks & Maes, MIT Press, pp.28-39.
[7] Soddu, C. and Colabella, E. (1997) *Argenic Design*, The European Academy of Design Contextual Design, Design in Context Conference, Stockholm.
[8] Spehar, B., Clifford, C.W.G., Newell, B. & Taylor, R.P. (2003). *Universal aesthetic of fractals*, *Computers & Graphics,* **27**, 813-820.
[9] Sudweeks, F. & Simoff, S.J. *Quantifying Beauty: An Information System for Evaluating Universal Aesthetics.* http://www.it.murdoch.edu.au/~sudweeks/papers/beauty.pdf
[10] Taylor, R. (2002) *Fractal Expressionism—Where Art Meets Science,* Art And Complexity, Ed. J. Casti, Elsevier Press.

[11] Thomas, D. (2003), *Aesthetic selection of morphogenetic art forms,* Kybernetes: The International Journal of Systems & Cybernetics, vol 32, pp. 144-155
[12] Todd, S., Latham, W. (1992). *Evolutionary Art and Computers,* London, San Diego CA. Academic Press.

APPENDIX A
Some examples of artwork created by CAT

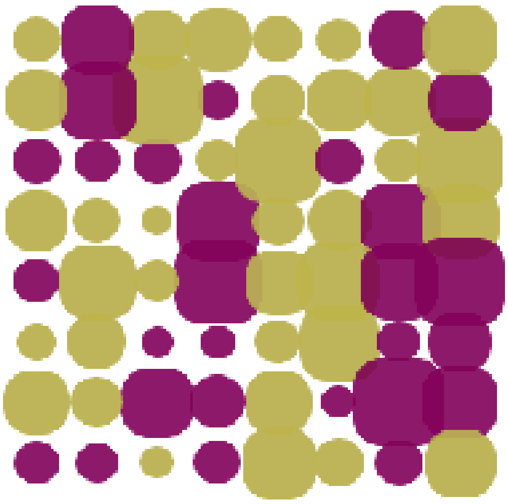
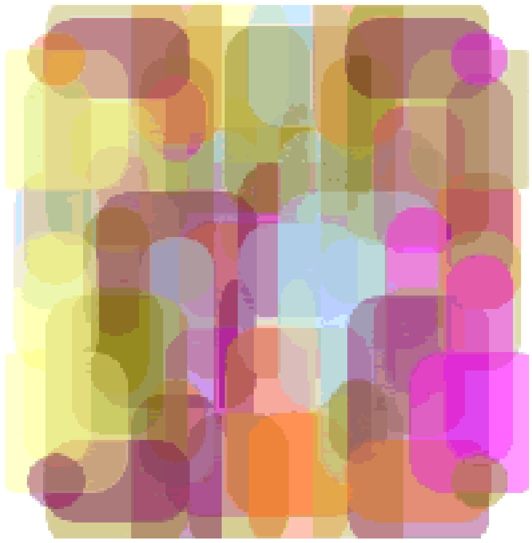
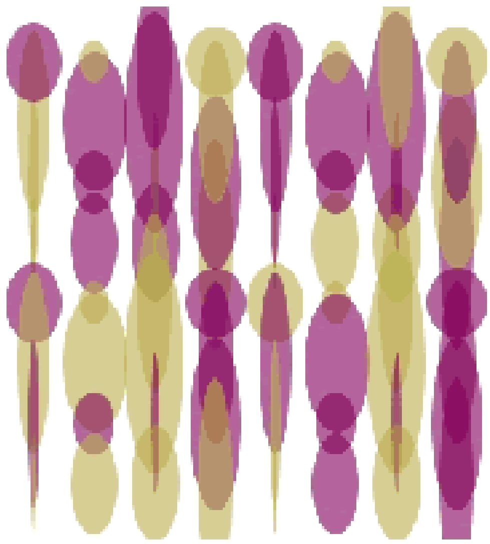
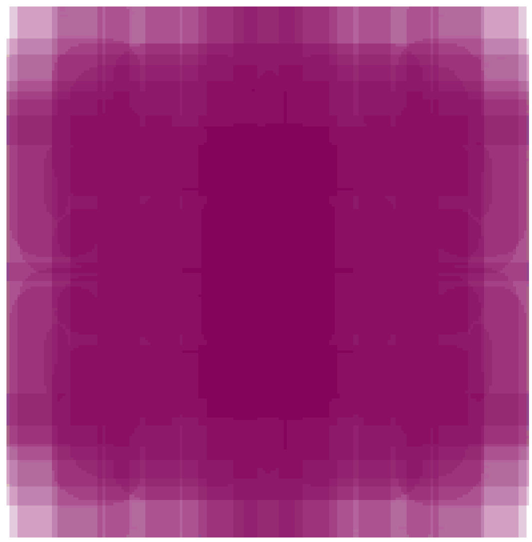
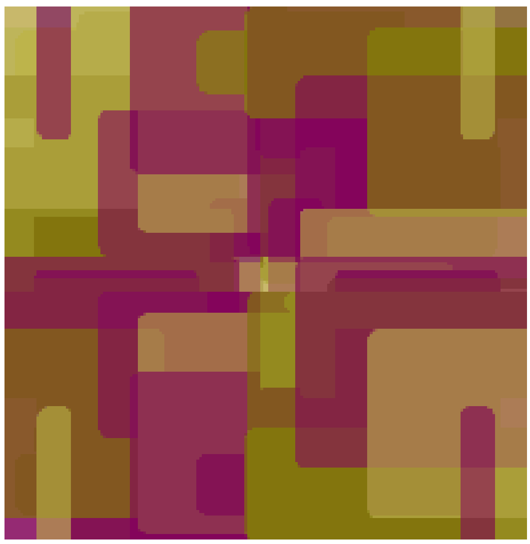
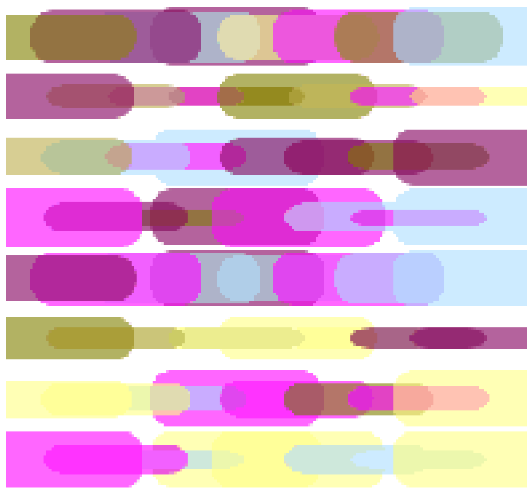


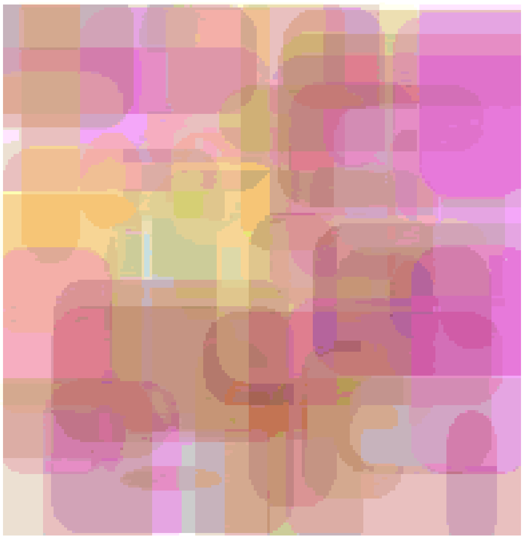
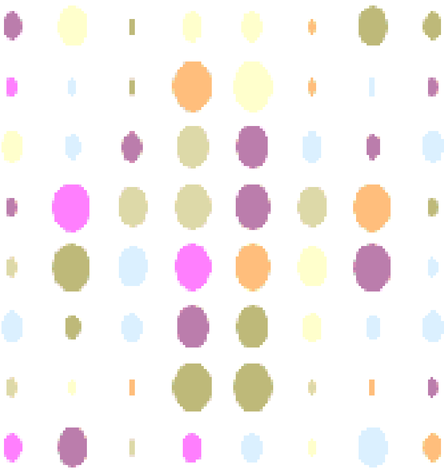
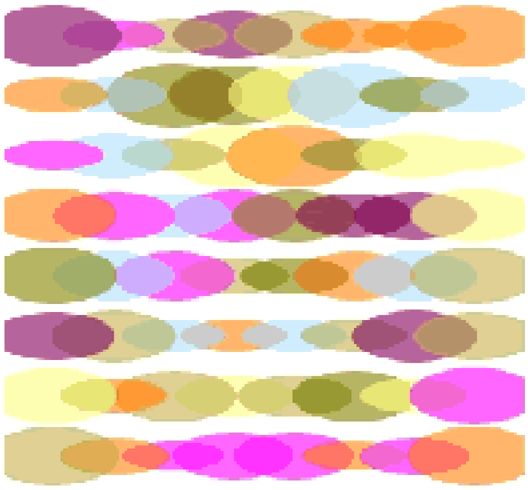

APPENDIX B

The Code
Java 1.4.2