School of Computer Science and Software Engineering
Monash University

Bachelor of Computer Science (Honours)
Clayton Campus

Thesis — Semester 2, 2004

Automatic Art Generation

Gregory Robert Santucci 12254266

Supervisors: Dr. Alan Dorin,
Dr. Jon McCormack

Contents

List of Figures

Abstract

This project intends to explore the automatic production of interesting images using computer programs. In particular, to discuss aesthetic theories, then to propose a system that expresses an artistic idea of what constitutes an interesting image, and finally to determine whether the system does produce interesting images exhibiting those artistic ideas. To limit the scope of this question, only abstract art is produced consisting of a limited set of primitives. The image generation framework discussed is a hierarchical, scene tree based method in order to guide the generation of coherent images.

I declare that this thesis is my own work and has not been submitted in any form for another degree or diploma at any university or other institute of tertiary education. Information derived from the published and unpublished work of others has been acknowledged in the text and a list of references is given.

Acknowledgments I would like to thank my supervisors Dr. Alan Dorin and Dr. Jon McCormack for their advice and insights into art and aesthetics which they provided throughout the year.

1 Introduction

1.1 Artificial Art

Producing artwork is traditionally a human activity. Some artists believe that trying to mechanise it or automate its production is a meaningless excercise since they consider art a form of human expression. There are others, instead who hold the view that producing works of art, or "aesthetic objects" is "as rational an activity as anything in the natural sciences". [3] With this idea, and with the advent of computer technology, Max Bense and Georg Nees ran the first exhibition of computer generated "algorithmic art" in 1965. Some artists reacted with hostility: where is the intuition? The inspiration? What could possibly be the message of these works? The term "artificial art" was coined by Max Bense to distinguish it clearly from human produced art, in a similar way that "artificial intelligence" is used to distinguish the disembodied and often brittle intelligence encoded in programs from the more fluid intelligence we are more accustomed to in our daily lives. Some artists such as Manfred Mohr used the computer to explore the generation of forms of a certain type, in his case, cubes. Other artists such as Roman Verostko, were interested in generating artwork that was purely the result of an algorithm he devised. Harold Cohen began as a painter, then wrote a program called "Aaron" which is an expert system encoding his own artistic style. An artistic movement had emerged involving the collaboration between computer scientists and artists, to explore the new possibilities of art creation that the computer can offer.

Using computers to generate art allows the exploration of abstract spaces in a free, fluid way that wasn't possible before. A lot of computer generated "artwork" that we are accustomed to seeing today consists of swirling vivid forms with complex structures. David Em working as the "artist in residence" for the Jet Propulsion Laboratory used the high end machines at the time to create works that were described as "startlingly original and vivid, having a dreamlike, almost nightmarish quality". [16] These images provide an entertaining dose of "eye candy", and I believe they are interesting in their own right because as you view them, you're mind tries to identify structures and substructures, a hierarchy of relationships. The mind is very adept at pattern recognition, it arguably evolved for this purpose, so anything that provides "practice" for this will be stimulating and enjoyable.

1.2 Aesthetic theories

We just identified a possible aesthetic theory: an object is aesthetically pleasing if it causes the mind to spend time trying to discern the web of relationships in it. However, according to this idea, if an image is a noisy random image, the mind may spend a huge amount of time trying to discern structure when there isn't any. In reality, the viewer becomes bored or disinterested when confronted with a noisy image, since our minds are accustomed to tuning out noise. So we can append the idea that there needs to be the right mix of "order" and "complexity" in the image for it to remain interesting. Too much complexity, and the mind will tire trying to make sense of the image, not receiving any rewarding stimuli to encourage it to keep going. Too ordered, and the mind deciphers the image too quickly. If there is the right mix, there should be constant "nuggets" of encouragement to keep the mind interested in deciphering the image, and there will always be plenty of structure left to explore.

Birkhoff's aesthetic measure uses this same idea. In 1933, he came up with a general formula relating "beauty" to "order" and "complexity", $M = O / C$ where M is the "aesthetic value", O is the order and C is the complexity. [1] I believe Birkhoff was on the right track relating order and complexity to the "aesthetic value" of an image. The neuroscientist

Vilayanur Ramachandran [18] theorises that “grouping” is an important aspect of an aesthetic piece since as we discover groups, the mind produces a sort of “a-ha!” feeling that encourages the search for more groups, and higher order structures. He suggests that this is a useful survival trait that has evolved to encourage the animal to remain motivated in pursuing a mate through fog, or to search the terrain for food. Humphrey [11] suggests a theory that is parallel to this: we enjoy aesthetically pleasing images for survival reasons, since aesthetically pleasing images teach us how to classify objects in the world around us. We will spend a long time looking at an image we find interesting because we are being challenged to try to integrate it into our model of the world, our classification system. For example, Humphrey states that in his experiments, monkeys prefer to look at abstract paintings more than pictures of appetising but familiar food. Humphrey’s aesthetic theory is that aesthetically pleasing objects should exhibit “likeness tempered with difference”, which supports the idea of a balance of order and complexity. Something too ordered or too expected, or something too noisy and complex will be ignored. Humphrey suggests that variations on what is already known, but not too far away are enjoyable too, perhaps because they gently extend the boundaries of the world view of the viewer.

The petals of a flower all have an identifiable shape, yet each one is unique. In viewing a flower for the first time, the mind being the great pattern recogniser would ponder over it and integrate this knowledge: flower consists of petals that are unique, yet identifiable as petals. This is why a flower is pleasing, our mind was able to perform its rule discovery to decipher it, and discern its structure, and thereby develop the model of the world. A hopeless jumble of complexity would baffle our mind, and we would not learn anything from it or enjoy it.

Another theory of what constitutes artwork is proposed by the philosopher Goodman. [7] An art object exhibits 3 symptoms: repleteness, expression and composition. Repleteness means that more physical characteristics of the work are relevant than if it isn’t being considered an artistic piece. For example, the quality of a line is relevant in a line drawing, but not a technical diagram. Expression means that the piece is expressing something that is not literally part of it: we say an image is fiery, but it is not hot to touch. Composition means that every component of the image makes a meaningful contribution to it. These ideas are difficult to quantify, but they can be interpreted to mean ultimately that an art object exhibits a balance between order and complexity.

Using systems of proportions has the effect of unifying a work. In ancient times, artists and architects in Egypt and Rome used systems of proportions to structure their objects. This resulted in a uniform, identifiable style, which according to Ramachandran, is a good idea since it results in less effort required by the mind to decipher it. There is a particular proportion that appears often in nature and artwork that emerges from the Fibonacci sequence. It is called the “golden ratio” and some psychologists have run tests to determine whether rectangles or objects proportioned according to it are more aesthetically pleasing than other objects. The human body is also approximately proportioned according to the “golden ratio”. My contention is that it is more aesthetically pleasing because we are so accustomed to seeing it, our brains may be wired to like it. The Fibonacci sequence can be constructed by the recursive relation $x_i = x_{i-1} + x_{i-2}$ where the first 2 terms$x_0$and $x_1$are both 1. The ratio of successive terms in the sequence approaches the “golden ratio”, which is the value $\frac{1+\sqrt{5}}{2}$ which is approximately 1.61803. One of the properties of this irrational number is that you can square it by adding 1 to it, so it is a solution to the quadratic $x(x-1) = 1$. See figures 1, 2 and 3 for examples of the golden ratio proportion. (The calipers pictured demonstrate the golden ratio).

People have often wondered about the idea of how much of the meaning of an object is just intrinsic to the object, and how much is culturally influenced, or influenced by the structure and behaviour of our minds. Hofstadter [9] proposes the experiment of sending

Figure 1: Egyptian tablet proportioned according to the golden ratio. (Reproduced from Schneider, see [23])

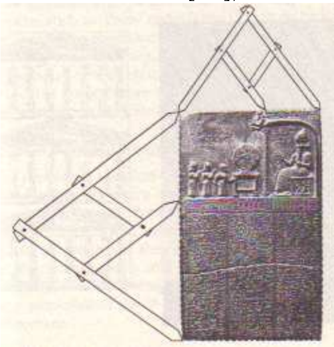

Figure 2: Notre Dame proportioned according to the golden ratio.(Reproduced from Schneider, see [23])

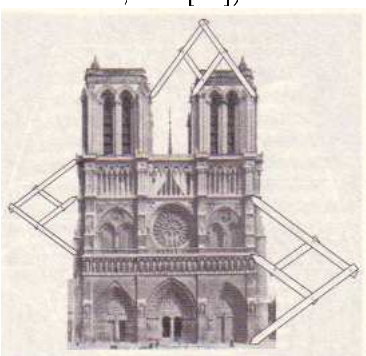

Figure 3: Golden ratio proportioning of the wrist and face. (Reproduced from Schneider, see [23])

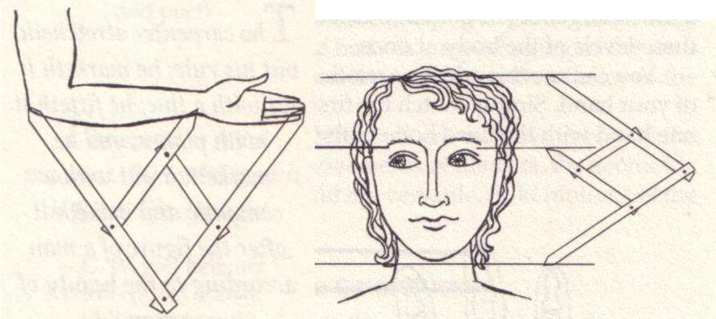

Figure 4: Left: The change in curvature of the logarithmic spiral is based on the fibonacci sequence, and logarithmic spirals are often seen in nature because they result in optimal packing. Right: The logarithmic spiral can be seen here being used as the basis for placement of scales on a pinecone. (Reproduced from Schneider, see [23])

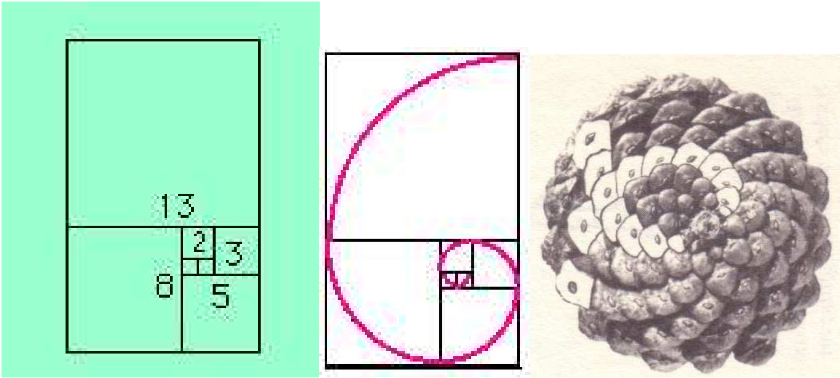

a record into space that plays the fugues of Bach, and seeing whether all the interesting structures and substructures, variations and developments in the music are discerned and bring as much joy to alien intelligences as they do to humans. Would alien intelligences also tend to prefer the "golden ratio"? If they do, would this be because nature tends to use it, because basing objects on the Fibonacci sequence is more economical? (See figure 4)

Certainly it is clear that what is considered interesting for a person varies based on their life experience, their education, and culture. If you have specialised knowledge, you may find some images interesting that to others would make no sense. For example, if you have prior knowledge of the way a particular artist structures their images, you'll be able to appreciate the levels of complexity in that artists paintings to a greater degree than for people without this prior knowledge. Again, based on Ramachandran's idea of stimulus and reward, an experienced viewer will be able to find the structures that elicit the "a-ha!" feeling whereas others won't. What remains consistent is that for a given viewer, an image which is too complex or too ordered would be considered uninteresting. The threshold is based on the person's prior knowledge.

1.3 Aim of the project

The aim of this project is to explore the automatic production of interesting images using computer programs. This is a very complex problem since what is considered interesting varies with culture, education, life experience and subjective opinion. One consistent feature of interesting images is that they exhibit a balance between order and complexity. I will explore the question of whether there is a set of artistic concepts that I can implement in a computer program to generate interesting artwork automatically.

A physicist in Adelaide named Dr. Alan Lee inadvertently shed some light on this question, when he was interested in finding out why the late work of Mondrian is considered so interesting and beautiful, when it seems to only consist of horizontal and vertical lines with some of the rectangles filled with a primary color. [17] The work of Mondrian seems so simple, arbitrary and abstract. (See figure 5)

He wrote a program to generate fake Mondrians. The program he wrote was intended to be the antithesis of the creative process. It would place the horizontal and vertical lines randomly, and color random rectangles. He ran a survey to see if art critics could tell the difference between the fakes and the reals, and found that they couldn't. So he concluded that they were not interesting after all. The professor was proceeding from the premise

Figure 5: Composition No. 10 by Piet Mondrian, reproduced from `www.ibilio.org/wm/paint/auth/mondrian`

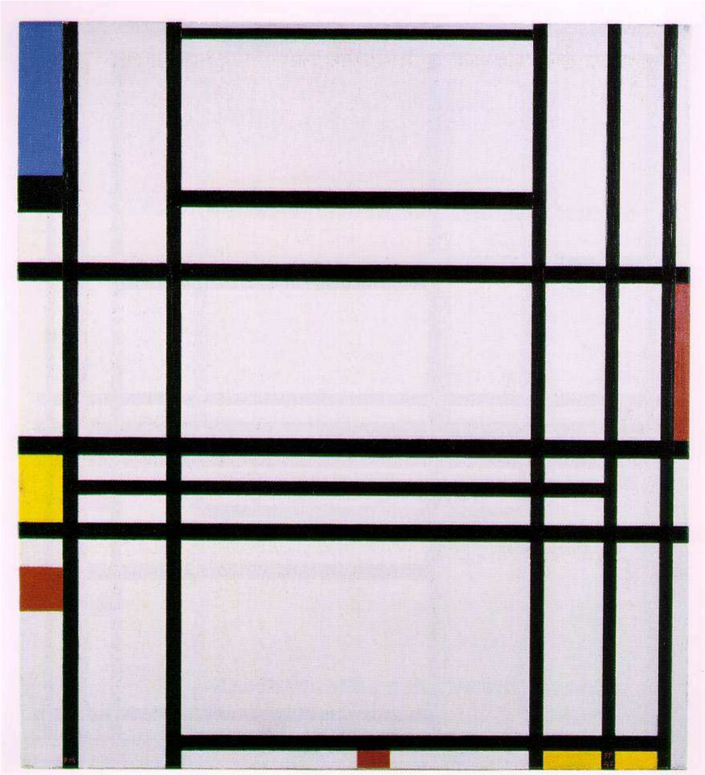

that to be interesting, an artwork had to be produced by a process that was interesting and complicated, that there had to be a "creative force" behind it. I believe the reason people find the work of Mondrian interesting is the fact that there are many ways of interpreting structures and substructures in his work, so there are many "a-ha!"s possible when viewing it. The fact that a simple process resulted in it does not have much to do with whether the image can be interesting or not, whether or not it can provide a source of continual stimulus for the mind.

Our minds enjoy trying to make sense of Mondrian images, however perhaps the motivation to do so may dissipate if we believe that there is no hidden message to decipher. This is the contention of psychologists who ran a study to determine if college students could determine if the aesthetic pleasure of an abstract painting varied if it was oriented incorrectly. The psychologists found that it doesn't. This could be explained by the lack of artistic knowledge of the students (they were not versed in the style of the artists in question).

In section 2, some psychological theories about interesting images is presented. In section 3, the way artists have written and used programs to create artwork and describe their artistic style is discussed. In section 4, some aesthetic measures based on the notion of a balance between order and complexity are presented. In section 5, some cognitivie models are described. In section 6, The work and philosophy of some abstract artists is briefly described. In section 7, the drawing program I wrote is introduced, and in section 8 the kinds of artistic effects it is capable of producing are revealed. In section 9 I summarise what the program has achieved in terms of being able to produce interesting images.

2 Psychology, cognition and interesting images

Our minds are adept pattern recognising machines. This is a trait that according to neuroscientists such as Ramachandran, evolved so we could process the information we need about our environment quickly. Aesthetic pleasure is something Humphrey says comes from things that our mind recognises can improve our internal model of the world. Therefore, aesthetically pleasing objects are recognised by our mind as things that can help us to learn.

Image processing occurs at many levels in the mind. In an early image processing stage, edge detection is performed, to identify boundaries of objects, and therefore, their forms. In later stages, basic groups of similar objects are identified. As this work is going on, the mind may request more information about an image, causing the eye to scan over certain parts of the image. Some researchers have noticed how the eye movements across an image varies according to the question the viewer is being required to answer about it. [24] This is consistent with the idea that the mind is trying to find some structure to the object subconsciously, so the mind can effectively extract the information it needs. Another experiment to illustrate the information processing mechanisms inside the mind is to present it with an image exhibiting what is known as "conflicting organisational patterns". As the mind explores one way to structure the image, another way to structure the image presents itself. (See figures 6, 7) Interestingly, the placement of scales on a pinecone also exhibits "conflicting organisational patterns" since when you discover one spiral, you can discover immediately another spiral going the other way that shares some of the same scales. The interesting thing these images have in common is that they all consist of a repeated, readily mechanisable pattern. Figure 6 is just a grid of squares, and figure 7 is just a tessellation. This suggests that the processes to generate interesting images may not be that complicated, it may not require a complex cognitive model.

2.1 Ramachandran

Ramachandran's main idea about art is that it is not about reproducing reality, but about deliberate hyperbole, exaggerations, and distortions to produce a pleasing result in the brain. [18] To support this, he proposes the idea of "peak shift". When you view an object, the part of the brain that can recognise it lights up and so you recognise it. If you accentuate some of the features of the object that distinguish it, you produce an image that can excite that "recognition" part of the brain even further. This is a technique cartoonists use to make their caricatures seem more real than the real thing, and it is why red herring gull chicks when feeding peck more vigorously at a yellow stick with three stripes than their actual mother. Ramachandran therefore proposes that the celebrated artists in our culture are successful because they have tapped into "figural primitives of our perceptual grammar". They intuitively know the space we think within, an identify features of that space to bring to our attention. Ramachandran cites the example of the Indian bronze statue "Parvathi". The artist has looked at the "posture space" of men and women, and is accentuating the difference, resulting in a statue of a woman that seems to be more feminine than a real woman.

2.2 Humphrey

Humphrey [11] tries to explain why humans should have evolved to enjoy art at all. He suggests the reason could be that new and interesting stimuli offer an opportunity for a creature to improve its internal classification system. Through aesthetically pleasing forms, a creature "learns how to learn". He describes his experiments with monkeys, where he found that the monkeys preferred to look at an abstract painting rather than a picture of

Figure 6: Conflicting organisational pattern consisting of squares. As soon as the mind discovers one cycle, another presents itself, using some of the same drawing elements. Reproduced from "Cognition in the Visual Arts" [24]

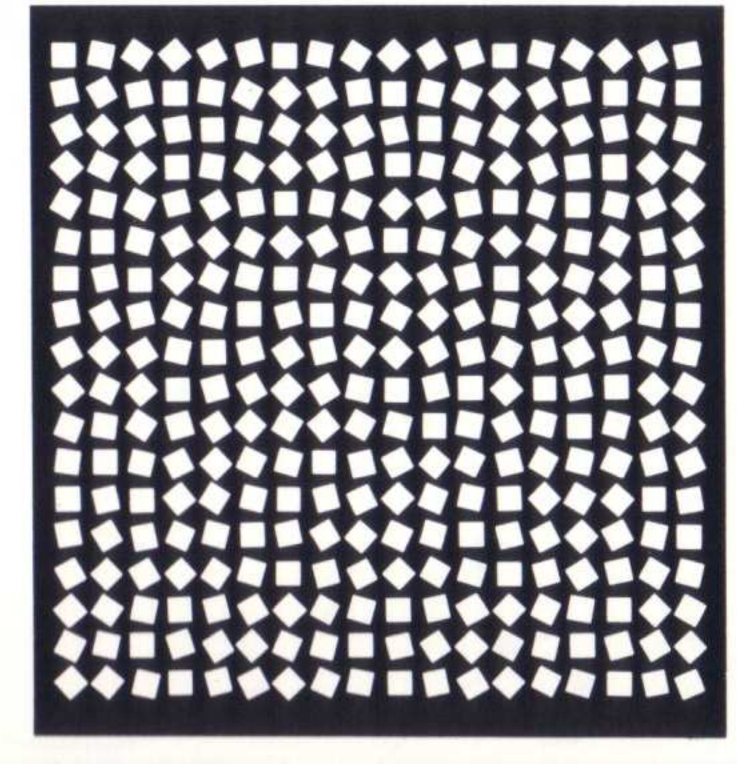

Figure 7: Conflicting organisational pattern. When a structure is discovered, another structure using some of the same components presents itself. Reproduced from "Cognition in the Visual Arts" [24]

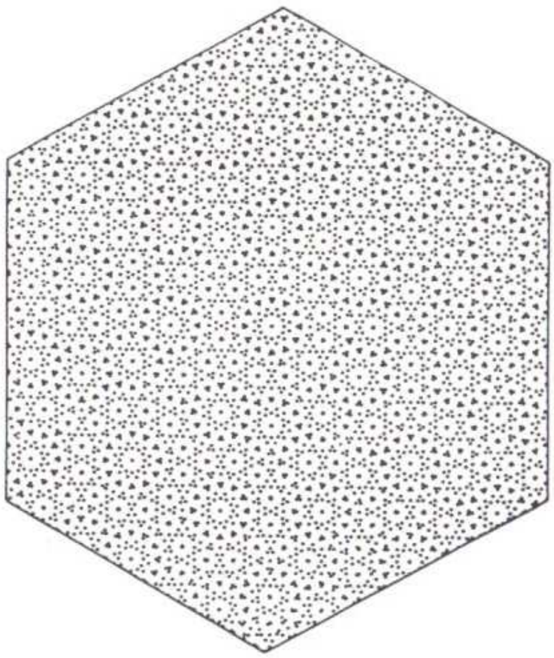

Figure 8: A drawing generated by a drawing robot controlled by the “Aaron” program

appetizing food. He concludes that the monkey's mind is drawn to an experience that can improve its view of the world, and therefore its ability to gather the information it needs, and therefore, its ability to survive.

2.3 Summary

Since interesting images are images that allow the mind to perform its pattern recognition and identification of structures and substructures fluidly, a program to generate interesting images should generate coherent forms with many details and relationships to explore.

3 Generating artwork

3.1 Aaron

“Aaron” [4] is a system which encodes rules and techniques for generating artwork. The author, Harold Cohen, spent 15 years developing this system, adding and refining the reportoire of rules and techniques that the program uses to generate art. Therefore, the program could be considered an expert system for the domain of generating art. However, as Cohen acknowledges, the program has nothing to say about creativity, or why the images are interesting. The program generates images of a distinct style.

Cohen notes that the methodology he uses relates to the modelling of expert systems since it relies heavily on his expert knowledge of image-making. [5] Cohen built Aaron with a motivation to discover something more about the art making processes than his work as an artist would allow. He feels he found in Aaron an "existence proof of the power of machines to do some of the things that we had assumed required thought, and creativity and self awareness of a human being".

When exhibited, the initial assumption of many people viewing the work was that the images must have been made in advance by the "real" artist. They could recognise his hand in the generated drawings. Once they hear the explanation, people start attributing intentionality to the program. Aaron seems to provide a convincing simulation of human performance. (See figure 8)

Aaron is designed to reflect Harold's understanding of what the human image-making process might be like. Its behaviour is to constantly shift its attention to various levels of detail, and that no single part knows what the drawing should turn out to be like. Creating the drawing is an iterative process: each placement of an object in the drawing constrains the subsequent placements. In each iteration, the "planning" component decides on a figure and asks the "mapping" component to provide space. However, the communication can happen the other way too. During iteration, the "mapping" component may announce the existence of a space, and the "planning" component will decide what to do on the basis of its availability.

3.2 Roman Verostko: The Algorists

This is an art movement [25] where artists use algorithms to create forms and compose images. They use the term "epigenetic" to describe algorithms which generate forms, claiming that the software is the genotype. This is because the software describes how to construct an image and what calculations need to be performed, so by executing the program, the image is created.

The artist built something like an expert system to control the behaviour of a plotter so it could control a paintbrush in a variety of ways. So although the computing power at the time was modest (1986), by using a plotter, the artist was free to try all sorts of experiments and ideas. In this way, the computer was used to run an image making process the artist devised to make a set of images. However, the computer wasn't used to do a lot of searching of a design space or evaluation, but only used to realise a technique.

Verostko describes the iterative process of developing the drawing software, and the delight he felt in this. The machine would realise a family of images, and he would modify the program rather than modifying the artwork to fine tune the images the program was generating.

Since the software can be written to "embody the procedures for artistic improvisation" and can therefore generate "innovative variation on the artists theme without the artist being present", Verostko claims that it opens the door to the question of authorship and originality.

Verostko describes the motivation for some components of the program in an interesting way. He states that if the program is to draw a line, then the:

> "basis for all the line drawing procedures must be identified. The basis for the starting point, color, quality and character of the line, changing angles, flow, length"

He is trying to describe the programming process that requires him to formalise concepts.

3.3 Braitenberg vehicles

Since simulations of physical systems and systems consisting of simple interacting components exhibit complex behaviour, these simulations have been used to form the basis of some computer generated artwork. By tracing the trajectories or other properties of some elements of the simulation, an interesting image may result. An example is tracing the trajectory of Braitenberg vehicles over a canvas.

Braitenberg vehicles are automata with sensors and motors. [2] With a simple wiring, the vehicles can exhibit object avoidance or object seeking behaviour. By setting up the group of vehicles to follow or evade each other and tracing their paths, a variety of images can be generated. These images end up consisting of a set of sweeping curves, which appear to be have some relationship to each other. The images vary based on the initial placements of the vehicles on the canvas, and which vehicles they have been wired to evade or follow. By initially placing the vehicles on the canvas in a particular way, images with many axes of symmetry may also be produced. Braitenberg vehicles can exhibit complex behaviour and complex trajectories through interaction with a complex environment. GroupC [8] have used Braitenberg vehicles to produce complex artwork, all with a clearly identifiable style. The images consist of winding complex hair like curves that appear have a relationship

Figure 9: The trajectories of simplified braitenberg vehicles initially placed randomly and wired to follow one another was used to produce these images.

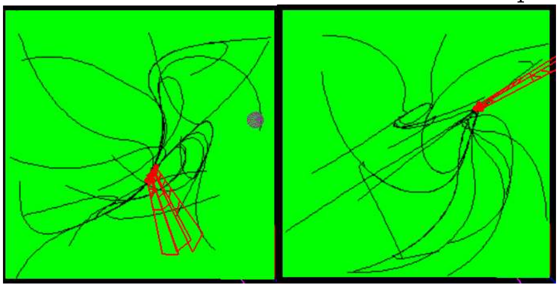

Figure 10: Image generate from the trajectories of Braitenberg vehicles (Reproduced from the GroupC website, `www.groupc.net`)

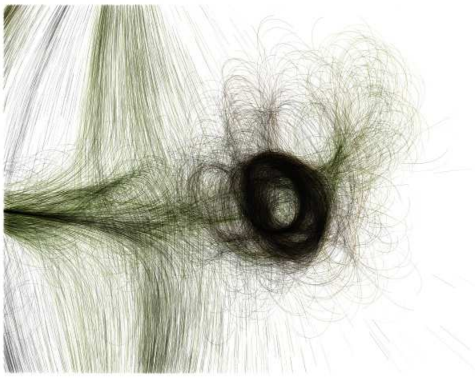

with one another. In figure 9 is a simple demonstration of this with a program I wrote to illustrate the principle, and in figure 10 is a demonstration with work done by GroupC.

3.4 Fractals

Fractals are images that are constructed with a motif which is repeated in a regular way at an ever diminishing scale. A fractal is interesting because the "a-ha!" feeling comes when it is noted that the motif is the important thing to recognise, since it unifies and describes the whole piece. However, upon subsequent viewings of other fractals, they become less interesting, since there is nothing new to explore once the motif is recognised. However, what remains interesting are the varying higher order forms that emerge when constructing the fractal from a motif. The mind tries to find a relationship between those higher order forms and the original motif.

When people see a fractal for the first time, they are often dazzled by the recursive nature of it, the intricacy, the fact that simply repeating a form in a regular way lowering the scale each time could produce an image that the mind can spend a while enjoying and finding many ways to structure. Fractals have a very hierarchical structure that can quickly be discerned by people, however, there are interesting relationships among the drawn elements at varying levels of the hierarchy that the mind enjoys exploring. See figure 11 for an example.

3.5 Summary

Possible approaches to generating computer art include building an expert system, writing a simple image generation program to generate a class of images, or basing the image generation on a simulation of a dynamical system. The approach I have taken described later

Figure 11: Fractals (Reproduced from Lauwerier [15])

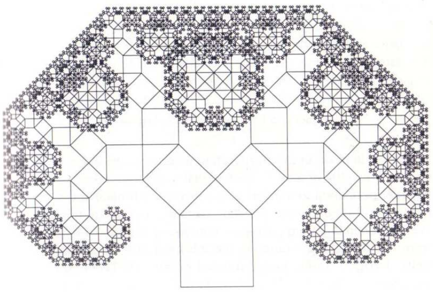

is to implement a general image generation program with a small reportoire of components to discover what kinds of artistic effects can be achieved with it.

4 Aesthetic measures

4.1 Birkhoff's aesthetic measure

Birkhoff's aesthetic measure [1] relates order and complexity to the aesthetic value of an object. The general formula Birkhoff states is M = O / C where M is the "aesthetic value", O is order and C is complexity. Birkhoff provides an interpretation of this formula for polygons and for vases. For a polygon, the complexity is the number of segments the polygon consists of and the order depends on the presence of vertical symmetry, point symmetry and mechanical stability with respect to an imaginary horizontal plane. Birkhoff demonstrated how his measure could be applied to vases as well.

By inspection, the formula always penalises complexity. This is contrary to the idea that an image should require the mind to explore in order to be interesting. Also, the formula always rewards order. In fact, according to the version of the formula for measuring the aesthetic value of polygons, the most pleasing polygon possible is a square. A square has many axes of symmetry, it is mechanically stable, and it has only four segments. The triangle fares worse, since although it has fewer segments, it has far fewer symmetry possibilities. (See figure 12)

Since this aesthetic measure is described in a formula, some psychologists have conducted controlled studies to statistically determine its validity. In a study involving people deciding the aesthetic value of polygons, Eyesenck [6] found that M = O / C does not provide a good fit to the data, and that M = O * C provides a better fit. This makes sense according to Birkhoff's definition of order and complexity for a polygon. A polygon with many sides, and which exhibits many different kinds of symmetry will be more interesting, than a polygon such as a square with very few sides and many different axes of symmetry since it is more intricate, there are more structures to explore. From the point of view of information theory, I think a polygon with many sides can only be arranged in very few ways so that it meets many symmetry constraints, so the information content of such an image will be quite high. Expressed another way, a polygon with many sides with the sides oriented randomly has a very low probability of exhibiting many symmetry constraints, however this probability increases as the number of sides decreases. This is perhaps the reason fractals appear interesting as well.

I think that relating order and complexity to the aesthetic value of the object is generally correct, since as neuroscience has found, the mind is an adept pattern recogniser, so it needs a stimulus it can play with. However, always penalising complexity and always rewarding order will not produce a good aesthetic measure. Perhaps instead, Birkhoff was describing a measure of the "orderliness" or "elegance" of the polygon.

Figure 12: Birkhoff's aesthetic measure applied to various polygons (reproduced from Scha and Bod [22])

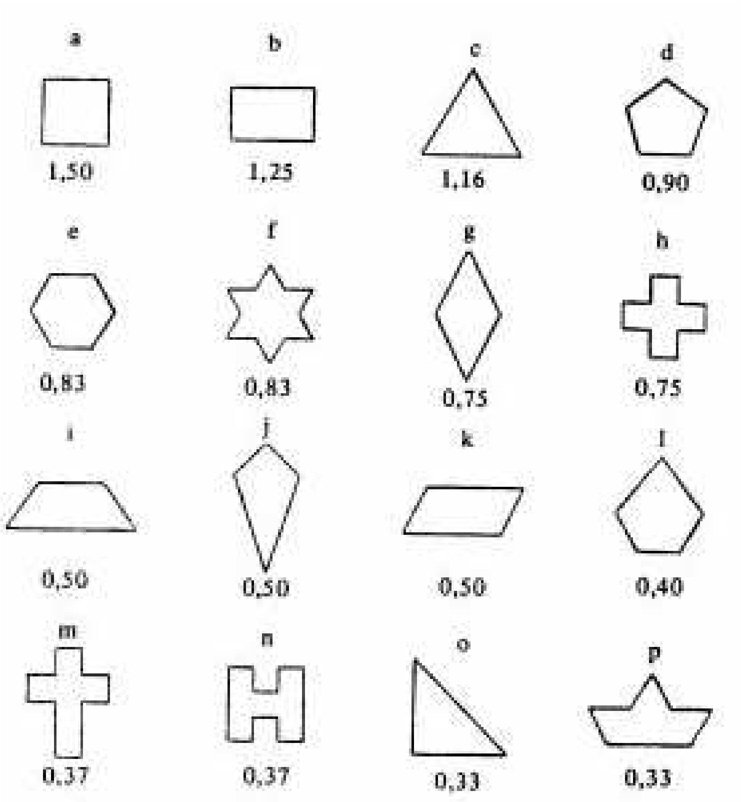

4.2 Information theory

The information theory method [14] based on Claude Shannon's information theory, tries to determine how much information there is in an image. It intuitively seems useful to measure the information content of an image, since this may suggest how much there is to explore in an image, and therefore, how interesting it is. It is naturally applicable to images that are composed of a grid of pre-defined symbols. Klinger and Salingaros [14] have invented a simple pattern measure based on two components:

1. Level of interest

   The level of interest is the measure of how many different symbols there are. This includes symbols that emerge from different partitionings of the image into square shaped regions.

2. Complexity

   The complexity is the measure of how many symmetries there are. The image is tested against a discrete set of axes and rotational symmetry angles.

The formula used is L = TH, where L is the "pattern measure", T is the "level of interest", or the number of discernible units, and H is the "complexity", or number of symmetries. According to Klinger and Salingaros, this measure tries to distinguish between "complexity with organisation" and "disordered complexity or randomness". See figure 13 for results. These ideas and results are consistent with the idea of interesting images being based on the idea of a balance of order and complexity.

Figure 13: Pattern measure example given by Klinger and Salingaros [14]. In decreasing order, the pattern measure results for these patterns were VI, V, II, III, IV, I (Reproduced from Klinger and Salingaros [14])

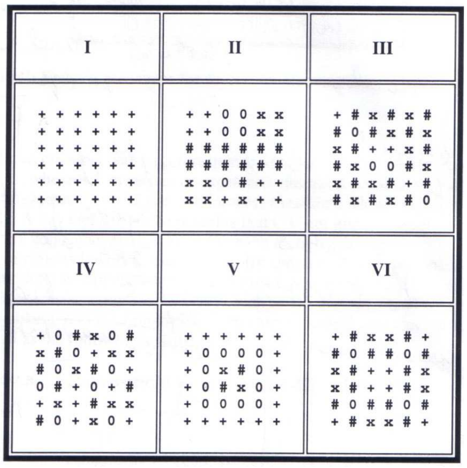

4.3 Summary

The aesthetic theory described here is based on the idea that an interesting image should exhibit a balance between order and complexity. This idea is related to information theory as well.

5 Creativity

We have presented the idea the the aesthetic value of an image is related to order and complexity, and justified it in terms of what is known about the mind being an adept pattern recogniser. What would be useful is a model of how the mind performs its pattern recognition when viewing images, which can then be tested and used to automate the production of interesting images and designs. I will present a cognitive model proposed by Hofstadter [10], then the creative design agent proposed by Saunders [21] in the following sections.

5.1 A cognitive model

Hofstadter describes a cognitive model. [10] It consists of three components he calls a slipnet, a workspace and a coderack.

The slipnet is a network where every concept in the domain is represented as a node, and arcs in the network have concepts associated with them as well. The nodes have an activation level that also affects the strength of associated arcs in other parts of the network, so that activation levels are able to "slip" between adjacent nodes. Active concept nodes modify the probability that certain classes of codelets will be added to the coderack. Codelets are small programs that run a process associated with some concepts in the workspace. This can involve adding elements or constraints to the workspace. The workspace is the area where the system creates its composition or proposed solution to a problem. The problem solving ability of this system emerges from the interactions of all the simple codelets modifying and adding to the workspace.

Hofstadter applied this model to the toy domain of inferring an analogy rule of letter sequence transformations. For example, if "abc" maps to "abd", what does "mrrjjj" map to? Most times the system comes up with the solution "mrrkkk" which is the application of the rule "replace last group with successor". Sometimes however, it comes up with the solution "mrrjjjj" which corresponds to the "deeper" rule "change length of last group by successor" since it notes that "mrrjjj" is made up of a group of 1, 2 and 3 letters, and a, b, and c are the 1st, 2nd and 3rd letters of the alphabet being mapped onto the 1st, 2nd and 4th letters. The system has a temperature which affects how readily proposed structures in the workspace are removed and changed, so the lower the temperature at the end of the run, the more "satisfied" the system is with the solution. Hofstadter chose this toy domain precisely because it would facilitate describing and testing this cognitive model clearly. The results achieved appear to be similar to the results you would expect if you tried this analogy making test on human subjects, therefore, it seems like a plausible model.

It is interesting to consider this model in the domain of art generation and evaluation, because of how it is based on discovering groups and subgroups in the information the system is presented with. The system could be presented with an image, and try to discern its structure, just as the system tries to discern the structures and substructures in the letter strings when trying to form an analogy. If the system was able to discern structures, and then an overall structure, we can say that the system evaluated the model as interesting. If not, the image is too "complex" for the system to decipher. The system, through its slipnet, has a fixed set of concepts and relationships among them, therefore a fixed "view of the world". It may be interesting to study what images will be considered interesting by such a system, and what images will be considered interesting after a change to the slipnet.

5.2 A curious design agent

Seeking aesthetically pleasing objects that can help us learn is the act of being curious about the world.

Rob Saunders [19, 20, 21]discusses his invention of a curious design agent.

This is an interesting method of traversing a parameter space in search of a solution. The agent has a component which compares the behaviour of the system it is testing now with its behaviour in the past. Once familiar with the behaviour of the system in a certain part of the parameter space, the agent starts to get "bored", causing it to venture into a new part of parameter space. The agent will have an expectation of how the system should behave with the new parameters.

The curious design agent alternates between periods of "problem finding" and "problem solving". When "problem finding" the agent modifies the simulation attribute that defines the problem, something like an independent variable. Here the agent tries to find a situation that reduces the utility function substantially, "breaking" the solution. When "problem solving" the agent modifies the other simulation attributes to maximise the utility function. The agent uses a neural network called a "Self Organising Map" to encode its expectations of the problem space. When it finds something novel, new neurons are added, and it continues exploring that part of the problem or solution space until it gets "bored" again.

Saunders applied this to some constraint and design problems and indicated how it could be used as a designer's tool the designer could interact with naturally. As with the cognitive model proposed by Hofstadter, it may be possible to use this system to explore or play with an image, to see whether the system can discern structures in it or not, after the view of the world is encoded in the "Self Organising Map".

5.3 Summary

An accurate model for how the mind perceives images would assist greatly in inventing a system that could automatically create interesting images. Work done by Saunders to generate interesting designs automatically produces designs that the design agent considers

Figure 14: ”Contrasting Sounds” by Wassily Kandinsky, reproduced from www.ibilio.org/wm/paint/auth/kandinsky

interesting, but which are not considered interesting by humans. Some human artists have attempted to describe in detail the philosophy behind their work, and this may reveal some insights into how these cognitive models may be improved.

6 Artists

It could be argued that with the abstract art movement began with the invention of the photographic camera, since artists were now free to explore abstract spaces and abstract forms. Abstract artists sometimes wrote at length about the philosophy behind their work, and even described the formalisms that their work was based on. Interestingly though, these abstract works are enjoyed regardless of whether the viewer has read the philosophy and understood the message and intent of the artist, since the mind simply enjoys trying to make sense of these images. There are many relationships to explore among the drawn elements that can be explored independently of any deeper message the artist is trying to convey.

6.1 Kandinsky

Wassily Kandinsky, like Mondrian, was an artist who experimented with limiting himself to using a few basic elements in his artistic compositions. In "Point and Line to Plane" [13] he describes how he thinks they should be combined. He asserts an analogy between the visual field occupied by a painting and the temporal field occupied by music. He asserts that "there are precise laws of composition, and that they are the same in all the arts". For example, a point has a sound, and if it is left by itself with lots of empty space around it, its sound can be heard clearly, conversely, if there are objects near the point, the sound of the point combines with the sound of those objects, so the point is less clear. (See figure 14)

6.2 Mondrian

In the later stage of his career, Mondrian became interested in using purely neutral forms to express the pure forms he perceived that the universe is based on. In this way, he viewed the world in a similar way to Plato: Plato’s writings are based on the idea that there is some underlying, pure form that can be expressed that we can gain access to. In his book “Plastic Art”, Mondrian talks about the “direct creation of universal beauty”. (See figure 15 on the following page)

Figure 15: "Composition III" by Piet Mondrian reproduced from `www.ibilio.org/wm/paint/auth/mondrian`

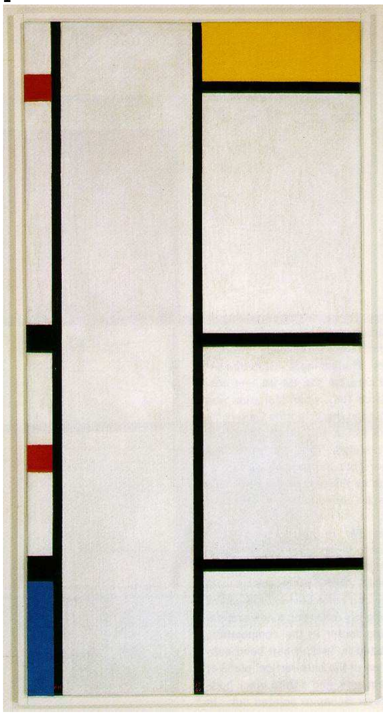

6.3 Johannes Itten

Itten [12] describes a set of contrasts that art objects exhibit to achieve various effects. They include transparent-opaque, smooth-rough, much-little, soft-hard, and light-heavy. When someone hears these terms, they are associated with experiences in day to day life, so the viewer may identify with and understand the structure and intent of images exhibiting these. To convert these ideas into something implementable in a program, it is useful to think of them in terms of changes in texture, in pattern, rhythm, and density.

6.4 Summary

As I demonstrate in the next section, I implemented a program which can exhibit many of these qualities Itten identified through variations in texture, rhythm, pattern and density.

Figure 16: Example scene hierarchy

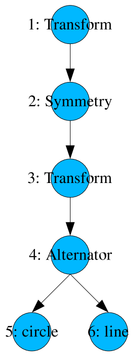

7 Method

7.1 Scene Hierarchy system

I implemented a system which maintains the state of a “drawing turtle” which moves across the canvas making marks. The state of the turtle contains the position, orientation, and axis scalings. The “drawing turtle” state therefore represents the current view transform, the current coordinate system.

To control the turtle, a scene tree system was implemented. At each node in the tree, a state of the turtle is maintained. This is so when the traversal of a subtree is complete, the state of the turtle is restored to that specified by the parent of the subtree. A node may modify the state of the transform (with a translation, rotation or scaling), draw to the canvas, and add its children to the stack any number of times, with varying transform each time. A variety of possible scene nodes can be implemented, and therefore, a variety of types of effects that can be obtained. For example, to implement a simple binary symmetry node its behaviour need only be to add all its children to the stack, flip the X axis, then add all its children to the stack again. (When a node is added to the stack, the coordinate system is saved with it). I decided on the scene tree system since a tree is a useful abstraction to use, since it is an intuitive, compact, expressive way to describe how to construct an image.

7.1.1 Example tree

The following diagrams show a tracing of the construction of an image from a scene hierarchy described in figure 16. In the following diagrams, the green line represents the X axis of the drawing turtle, the red line represents the Y axis of the drawing turtle.

Start stack: 1

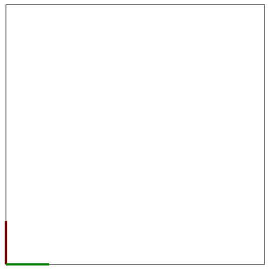

1. Popped node 1, move the turtle, push node 2

   Stack: 2

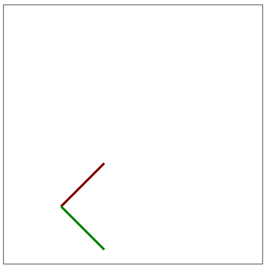

2. Popped node 2, add child node 3 to the stack, flip the X axis, add child node 3 to the stack, push node 3

   Stack: 3 / 3

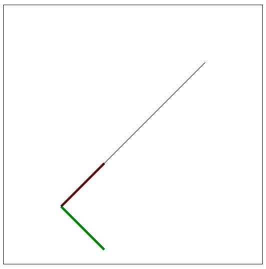

3. Popped node 3, move the turtle, push node 4

   Stack: 4 / 3

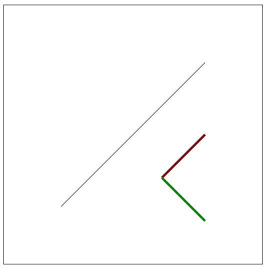

4. Popped node 4, alternator node, add child node 5 to the stack

   Stack: 5 / 3

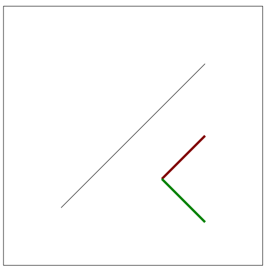

5. Popped node 5, circle node, draw a circle

   Stack: 3

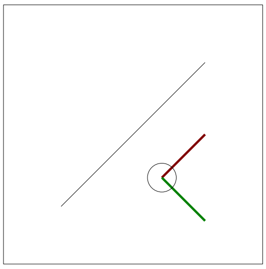

6. Popped node 3, move the turtle, push node 4 (Note how the turtle moves in the opposite direction, since the X axis has been inverted as a result of the transform saved to the stack in step 2)

   Stack: 4

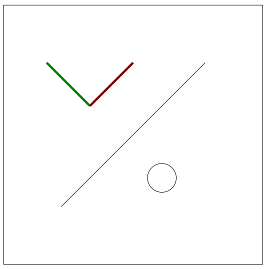

7. Popped node 4, alternator node, add child node 6 to the stack

   Stack: 6

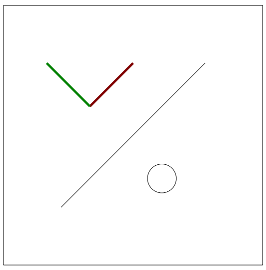

8. Popped node 6, line node, draw a line

   Stack:

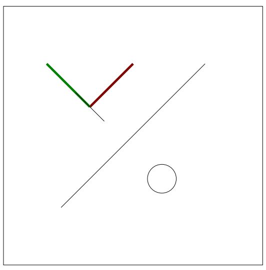

9. END

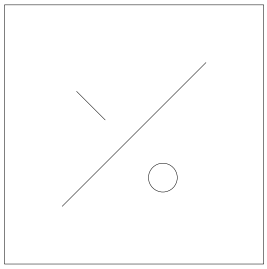

7.1.2 Node types

The following node types were implemented.

- *Symmetry node*

  The symmetry node will place all its children on the stack, then scale the X axis of the drawing turtle by -1, then place all the children on the stack again.

- *Radial symmetry node*

  The radial symmetry node will place all its children on the stack a certain number of times, each time rotating the turtle by a constant amount.

- *Transform node*

  The transform node will change the current transform matrix. This could be just a movement, or a rotation, or a scaling.

- *Alternator node*

  The alternator will place one of its children on the stack. Each time the alternator node is executed, it will place each of its children on the stack in turn.

- *Draw line node*

  The draw line node will draw a line 10 units long along the Y axis in the current coordinate system.

- *Draw circle node*

  The draw circle node will draw a circle of radius 10 units centred on the origin in the current coordinate system.

7.2 Functions to generate scene trees

To achieve multiple levels of radial symmetry, a branchless tree can be constructed that consists of radial symmetry nodes interleaved with transform nodes, ending with a draw line node.

Figure 17: Fractal images generated by the program using binary symmetry when the transform nodes are all the same.

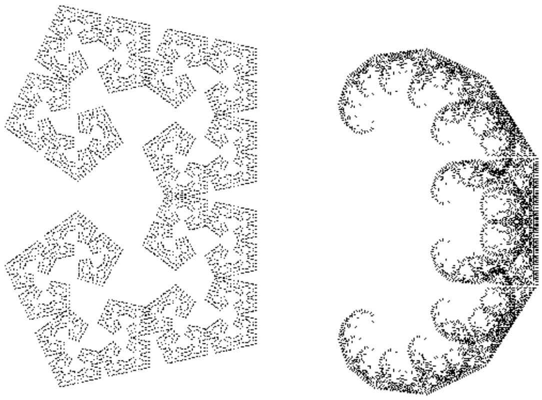

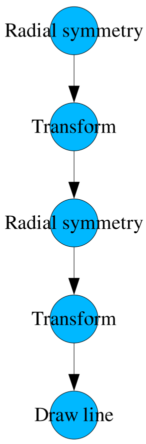

Upon each visiting of a radial symmetry node, the child is added to the stack multiple times with a rotation applied each time. The child node would typically be a transform node which moves the turtle forward and may scale the coordinate system down, then the next child node would be either another radial symmetry node or a draw line node. In this way, an image consisting of cycles upon cycles upon cycles of lines would be generated. The overall picture will have a circular structure.

To achieve multiple levels of binary symmetry, simply substitute all the radial symmetry nodes with binary symmetry nodes in the above diagram. Upon each visiting of a symmetry node, the child node is added to the stack twice, once with the X axis inverted. The overall picture will end up consisting of 2 distinct halves. The kinds of effects possible with a tree based on binary symmetry nodes are described in the next section.

To generate an image that looks like a fractal, I simply set the transform nodes to all be the same. This produces an image exhibiting self similarity. (See figure 17)

I was interested in images that could exhibit other artistic properties, so I set the parameters of the transforms randomly and independent of each other.

8 Results

To generate the following images, the parameters of the radial symmetry nodes (the stepping angle and the number of iterations) and the parameters of the transform nodes interleaving them (translation, rotation, scaling) were random within some range. The images all demonstrate coherent structures with plenty of relationships, patterns and substructures to explore.

8.1 Radial symmetry node based tree

By modifying the stepping angle of each of the radial symmetry nodes and the transforms, a very large variety of images exhibiting a variety of artistic features may be generated.

Figure 18: Radial symmetry star example

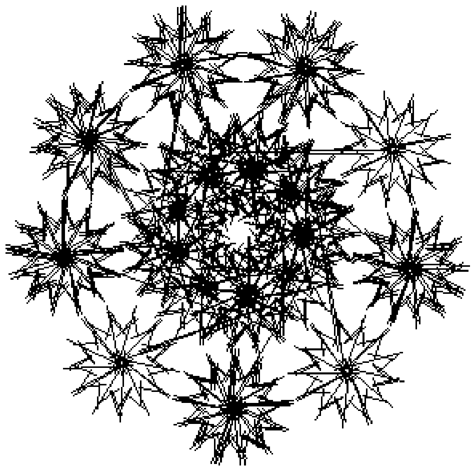

Figure 18 was produced with a tree with 3 levels of radial symmetry. Some emergent forms are apparent. The star shaped polygons around the outside for example. This image does not exhibit "conflicting organisational patterns" since once the main structures are discerned, such as the large star in the middle surrounded by smaller stars, there aren't other interpretations that are immediately apparent. The image is interesting since there are discrete structures that have emerged from the placement of the lines. Also, the quality of the lines varies throughout the image. In some places, the lines are bolder and more reinforced. This is achieved because the number of iterations of the radial symmetry node may vary, and so sometimes complete cycles are not completed.

Figure 19: Radial symmetry: emergent circles

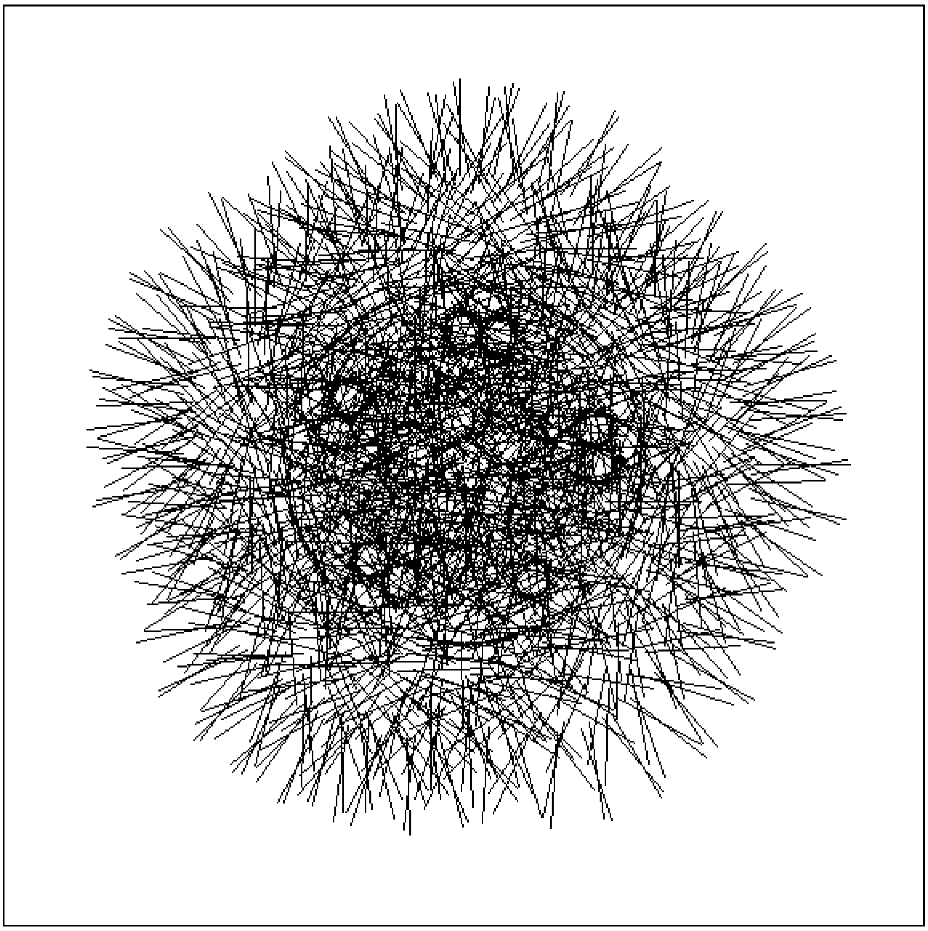

Figure 19 image exhibits emergent forms. In particular, outlines of circles can be seen toward the center. Moving further in, the image starts to appear noisy. Further out from the centre, more possibilities for discovering rhythms of the lines exist. The centre of the image is more dense than the outer part of the image, demonstrating the ability of this system to generate images with varying densities.

Figure 20: Radial symmetry: Rhythms

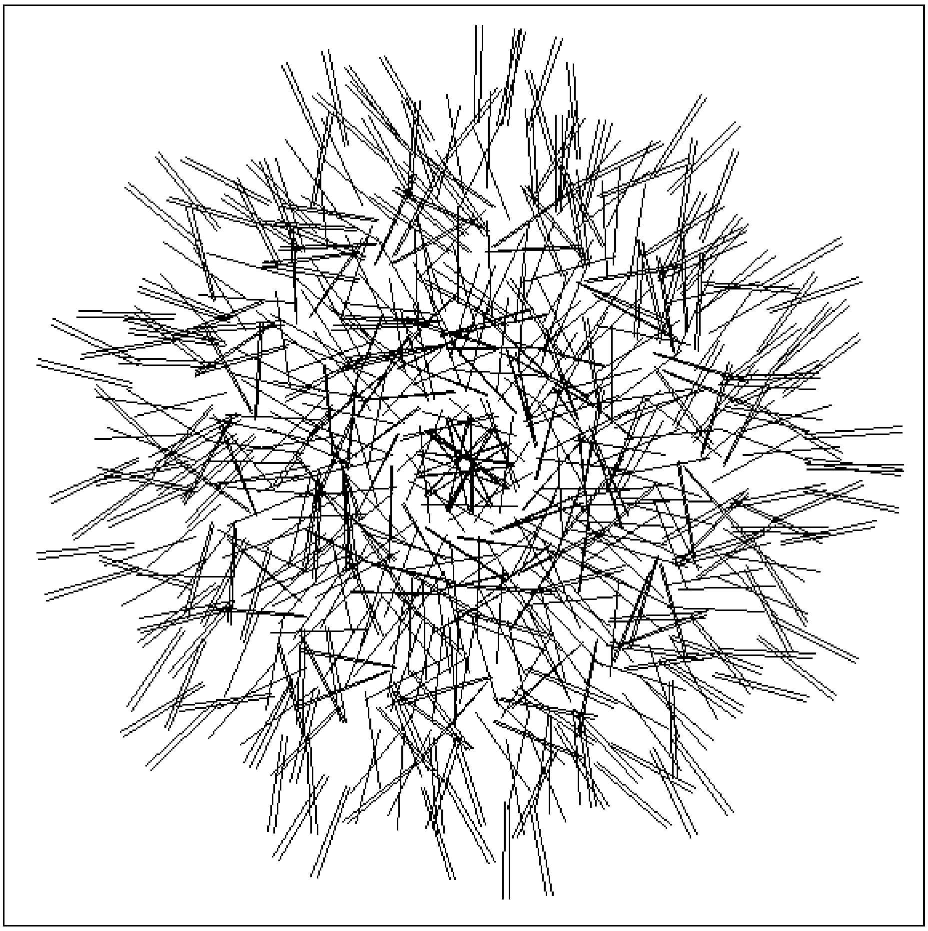

Figure 20 seems to have many rhythms and also seems to partially exhibit conflicting organisational patterns, since as you find one way of structuring the groups of lines, another way presents itself.

Figure 21: Radial symmetry: sparse

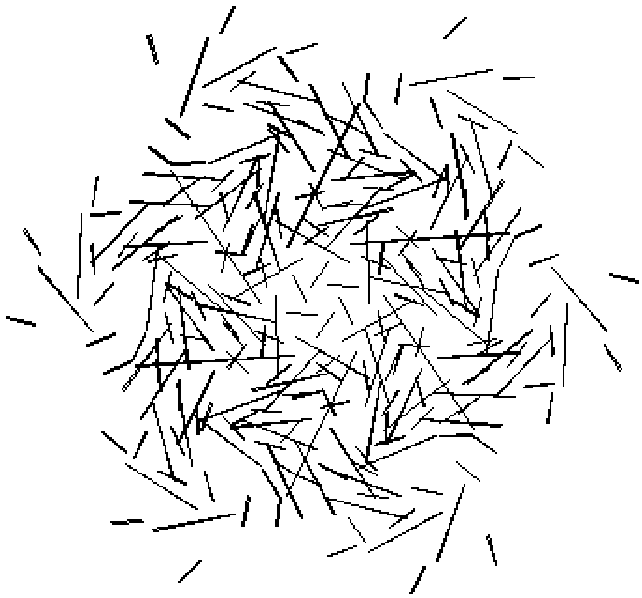

Figure 21 seems to be ordered, but defies any sort of discovery of consistent rhythms, so in this way it exhibits the idea of "likeness tempered with difference" described by Humphrey. There are collections of smaller lines around the edge of the picture.

Figure 22: Radial symmetry: regular

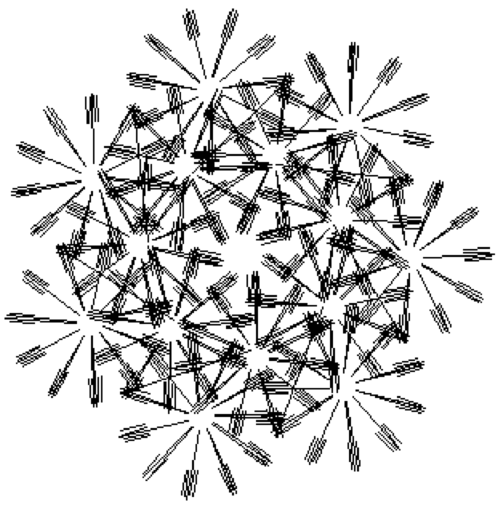

Firgure 22 demonstrates that by superimposing many copies of a pattern in this regular, circular way, other emergent forms become evident through the intersections of the lines and the figure formed in the centre.

8.2 Binary symmetry node based tree

The following images (figures 23 and 24) demonstrate that the system is capable of generating images with implicitly defined contours and regions of varying intensity. Also, the images have a recognisable, coherent structure. Every component and substructure seems to have a definite relationship with the whole, and so it meets part the criteria of "composition" that Goodman states and art object should demonstrate.

Figure 23: Binary symmetry: coherent 1

An interesting thing to note in figure 23 is the effect of a closed body with components and a definite form and direction. The texture varies across the form, and upon close inspection, there are always new structures and relationships to explore.

Figure 24: Binary symmetry: coherent 2

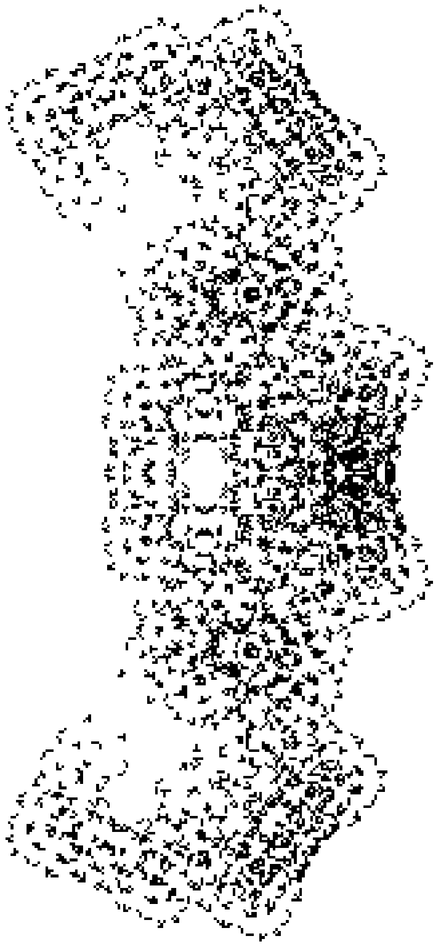

Figure 24 exhibits a more uniform texture, with a thinning out of the texture before the border where a definite contour can be perceived.

Figure 25: Binary symmetry: self similar with fur texture

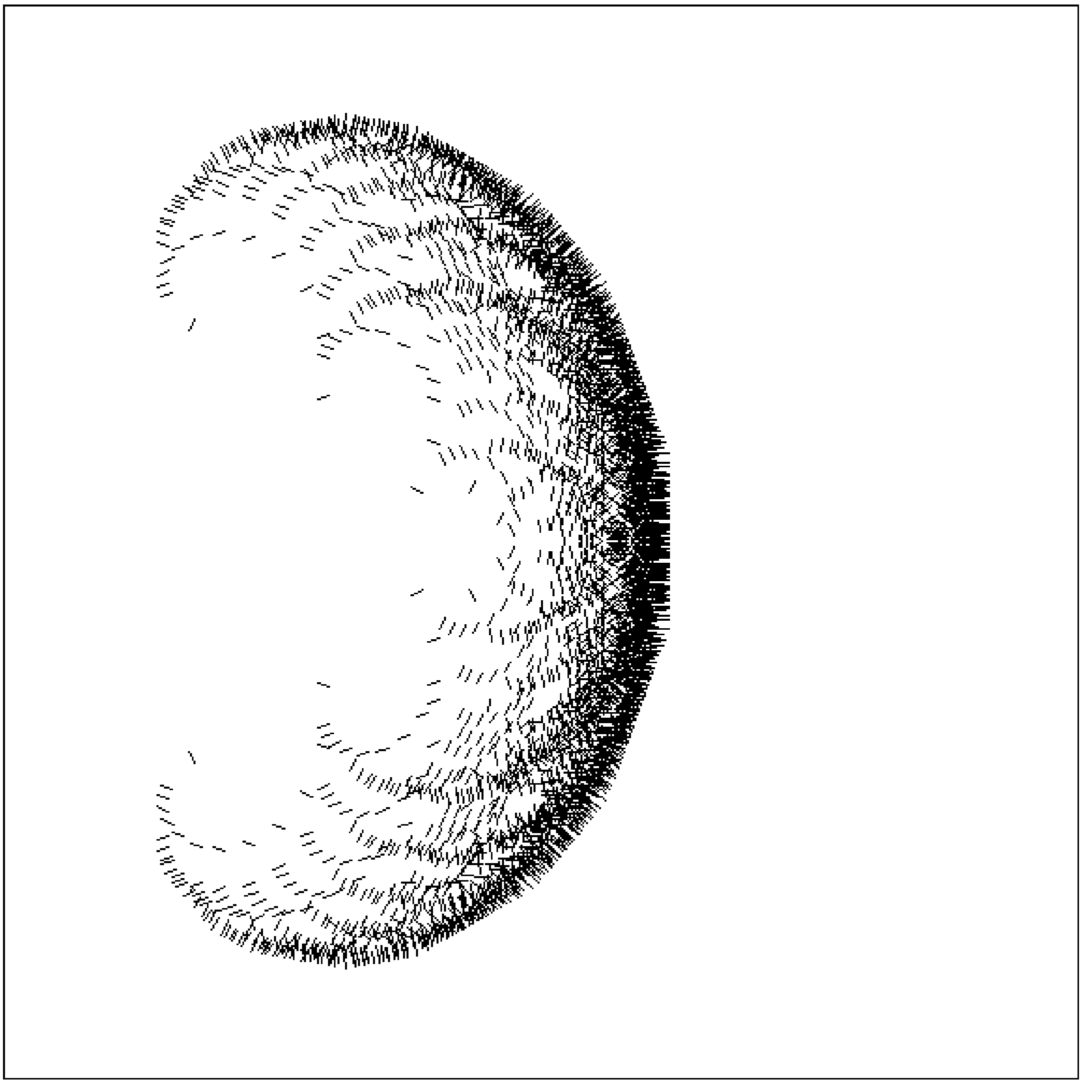

The tree for figure 25 consisted of 9 levels of binary symmetry. The system has generated a texture with greater density towards the edge of the form. Interestingly, this image is a self similar image since all the transforms in the scene hierarchy were the same when it was created.

Figure 26: Binary symmetry: contours and regions

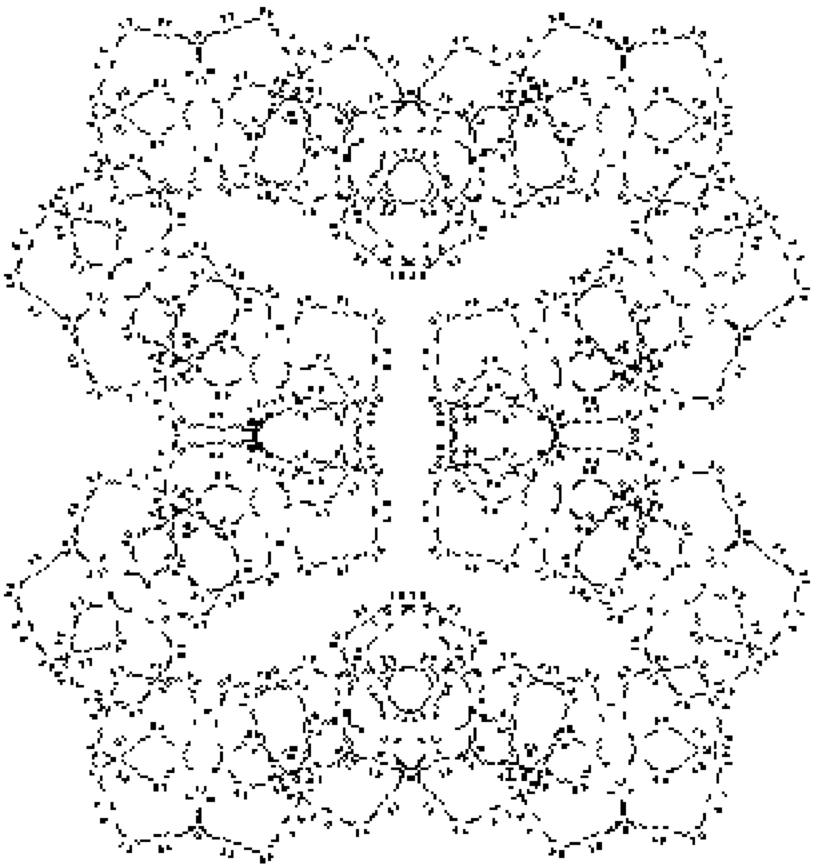

This image in figure 26 demonstrates the ability of the program to generate images with contours and varying shaped regions. The regions are not random, and they are not completely uniform: instead, they have a definite overall structure, and each region seems to have its own function and relationship with the whole. In this way, the piece reminds of the structure of living organisms. The tree for this image consisted of 9 levels of binary symmetry.

9 Conclusion

9.1 What was achieved

The objective I wanted my drawing software to achieve was to draw abstract, aesthetically pleasing images which exhibit qualities artists such as Itten have identified are important. Through the program I implemented, I discovered that an effective way to achieve this is to use many levels of radial or binary symmetry.

9.2 Conclusions drawn from the project and the ideas it investigated

Using many levels of radial symmetry or binary symmetry, it is possible to generate pictures exhibiting a variety of patterns, forms and textures. These images are interesting because they are structured, and there is plenty to explore, and many of the images exhibit "conflicting organisational patterns".

9.3 Possibilities for further research

It would be useful to find a way to specifically request a kind of texture or form desired, and have the program decide what the parameters for the radial symmetry nodes should be.

The program only generates images with a predictable overall structure. When using radial symmetry nodes, the overall structure is always circular, when using binary symmetry nodes, the overall structure is always 2 mirrored halves. New kinds of nodes based on these symmetry nodes that place children on the stack many times with different ways of varying the transform each time could produce new kinds of effects and image structures. Also, I have not looked at varying the structure of the tree very much, or explored the effects possible with the alternator node.

9.4 How what I did could be extended to deal with these possibilities

A simple modification to the program could allow the manual exploration of the space of possible settings of the parameters of the transform and symmetry nodes in the tree. The "Blind Watchmaker" method could be used so that the user could see the effects of many simultaneous parameter settings changes. In this way, at least the user could head towards a desired effect, even if it isn't clear how to quantify the effect.

References

[1] George David Birkhoff. *Aesthetic Measure.* Harvard University Press, Cambridge, 1933.

[2] Valentino Braitenberg. *Vehicles: Experiments in Synthetic Psychology.* The MIT Press, London, 1984.

[3] Candy and Edmonds. *Explorations in Art and Technology.* Springer Verlag, London, 2002.

[4] Harold Cohen. The further exploits of aaron, painter. http://crca.ucsd.edu/~hcohen/cohenpdf/furtherexploits.pdf. Website last accessed: 15th October 2004.

[5] Harold Cohen. Parallel to perception: Some notes on the problem of machine-generated art. *Computer Studies*, 4:3–4, 1973.

[6] W. Ray Crozier and Antony J. Chapman. The perception of art: The cognitive approach and its context. *Cognitive Processes in the Perception of Art*, pages 3–23, 1984.

[7] Nelson Goodman. *Languages of Art.* Hackett, Indianapolis, 1976.

[8] GroupC. http://www.groupc.net. Website last accessed: 15th October 2004.

[9] Douglas Hofstadter. *Godel Escher Bach: An Eternal Golden Braid.* Penguin books, London, 1979.

[10] Douglas Hofstadter. *Fluid Concepts and Creative Analogies.* Penguin books, London, 1995.

[11] N. K. Humphrey. The illusion of beauty. *Perception*, 2:429–439, 1973.

[12] Johannes Itten. *Design and Form.* Van Nostrand Reinhold Company, Ontario, 1963.

[13] Wassily Kandinsky. *Point and Line to Plane.* Dover Publications, New York, 1926.

[14] Allen Klinger and Nikos S. Salingaros. A pattern measure. *Environment and Planning B: Planning and Design*, 27:537–547, 2000.

[15] Hans Lauwerier. *Fractals, images of chaos.* Penguin, 1987.

[16] Donald Michie and Rory Johnson. *The Creative Computer.* Viking, Harmondsworth, 1984.

[17] Radio National. Radio broadcast transcript of the program Ockham's Razor. Is the artistic value of mondrian's art testable? http://www.abc.net.au/rn/science/ockham/stories/s645908.htm. Website last accessed: 15th October 2004.

[18] Valayanur Ramachandran. *The Emerging Mind.* Profile Books, London, 2003.

[19] Rob Saunders. Artificial creativity: A synthetic approach to the study of creative behaviour. http://www.arch.usyd.edu.au/~rob/study/publications/ccmcd01/SaundersGer%o2001CCMCD.html. Website last accessed: 15th October 2004.

[20] Rob Saunders. Artificial creativity: Emergent notions of creativity in artificial societies of curious agents. http://www.arch.usyd.edu.au/~rob/study/publications/second_iteration/Sa%undersGero2001SecondIteration.html. Website last accessed: 15th October 2004.

[21] Rob Saunders. A curious design agent: A computational model of novelty-seeking behaviour in design. `http://www.arch.usyd.edu.au/~rob/study/publications/caadria01/2001Saund%ersGeroCAADRIA.html`. Website last accessed: 15th October 2004.

[22] Remko Scha and Rens Bod. Computational esthetics. *Informatie en Informatiebeleid*, 11:54–63, 1993.

[23] Michael S. Schneider. *A beginners guide to constructing the universe*. Harper Perennial, New York, 1994.

[24] Robert Solso. *Cognition and the Visual Arts*. MIT Press, London, 1994.

[25] Roman Verostko. Software as genotype, a new dimension of art. *Leonardo*, 23:17–23, 1990.

A How the deliverables produced differs from those spelt out in the research proposal

In the research proposal, the direction of the project was to implement a genetic algorithm to evolve a drawing automaton (the control system for a drawing turtle) that would draw pictures which maximised some aesthetic measure. I had no knowledge about art and aesthetics, so I spent a long time studying aesthetic measures, and learning about art. I needed a theoretical basis for my work, so that it would not seem arbitrary and directionless.

After identifying an aesthetic theory to base my work on, I built a program that could construct images using a scene tree idea that would realise images based on this aesthetic idea. The idea was to implement new kinds of nodes, and look at what kinds of artistic effects they were capable of producing.

Also, I wrote a program which generates images based on the motion of a simple type of Braitenberg vehicle. I do not discuss that program at all in this thesis, however an example of the type of image it can generate can be seen in figure 9, near where I describe the work done by GroupC.

B Clarification of original contribution

My original contribution is my discussion of an aesthetic measure to do with the balance between "order" and "complexity", and how this is consistent with the idea that the mind tries to identify structures and substructures within an object. Also, I wrote a program that draws images exhibiting many levels of binary symmetry or radial symmetry, and demonstrated that this program is capable of producing interesting images since its images are coherent, with many structures and substructures, and the images demonstrate a variety of artistic effects.