Parallel to AARON

Aaron Steed 0311237

VT 3045 Fine Art Thesis

December 11th 2005

University of East London

School of Architecture and Visual Arts

BA Graphic Fine Art

Contents

Introduction

For over a year I have been concerned with the modelling of an artificial intelligence that will be able to create art. After a month of research a year ago I was introduced to the work of Harold Cohen who is the creator of AARON an artificial artist that has been in development for over thirty years.

To best understand the task I would aim to achieve I took it upon myself to use this dissertation as an opportunity to learn about the work of Cohen, artificial intelligence and technology that may be of use. This project's research necessitated creating models of the technology and theory I needed to understand better. As well as a comparison of theory I will be exhibiting some of the results of my practical investigations.

This thesis parallels Cohen's first paper on his work AARON that sought to propose what he would eventually create. Using Cohen's research papers, artificial intelligence theory and art theory I will outline how I myself may embark on a similar venture.

- Chapter 1 is on the subject of human perception and artistic perception as it is defined by theory and by the work of Cohen.
- Chapter 2 is on the subject of artificial intelligence. The various theories that currently exist on the subject and the technologies we have to investigate the subject.
- Chapter 3 looks at a number of digital artists whose work offers some parallel to Cohen's work that we might understand the nature of creating an autonomous creator of art.

- Chapter 4 is on the subject of artificial neural nets, which are modelled on the cells of the human brain. Research with computer models was conducted in this area to test their practical use.
- Chapter 5 concludes this dissertation with a proposal as to what would qualify as a system parallel to Cohen's work AARON.

The CD that accompanies this thesis has demonstrations of the neural net models I built to understand the different theories. Use of the CD is not essential to the reading of this thesis but is included as a digital appendix. The software requires the installation of Java on a PC using a Microsoft Windows operating system but requires no prior installation on an Apple Macintosh as Java comes bundled with the Apple Mac operating system. An executable that will install Java on Windows operating system is included on the CD. The CD is best navigated with an internet browser by opening the file named index.html. The various internet references in this dissertation have been put on the index page for the convenience of the reader.

A drawing made by an AARON application
(Application available from http://www.kurzweilcyberart.com/)

Chapter 1: Methods of Perception

An inquiry into the nature of art and human perception is valuable when proposing the creation of any analogue of human endeavour. I do not presume to copy the model of Mother Nature. But seeing as the product of the proposed machine is being created to satisfy one of mankind's pursuits it falls to us to understand that pursuit better. “Art is valuable to human beings by virtue of being made by other human beings, and the question of finding more efficient modes than those which characterise human performance simply does not arise.” (Cohen Harold, 1979, 1). He later remarks that he sees no compelling reason to model the human visual data processing method as a camera has little in common with the human visual system (Cohen Harold, 1982). When reviewing Cohen's work in comparison to visual theories though we find confirmation of his methods. Although AARON does not execute standard human visual perception - we will find that it has been built (unknowingly or not) with a great deal in common with its flesh counterparts.

The form of perception considered will specifically be vision, as it applies to the machine that is proposed. Although work on AARON is currently concerned with colour, it is simpler to start work on systems in black and white. There are two domains of the human visual system that Livingstone refers to as the What system and the Where system. The Where system deals with 'luminance'. This is perceived lightness as opposed to a quantity of photons hitting a meter. It is the rod cells in the human eye that tell us of luminance, giving one the ability to perceive spatial depth, organisation and movement. The registering of these aspects of vision is insensitive to colour. The What system, dealing with colour is by comparison slower and less sensitive (Livingstone Margaret, 2002). Being concerned with the identification of objects it is of course of importance to Cohen in his current research. AARON needs to effectively deal with colour now that its work has progressed to a representation stage.

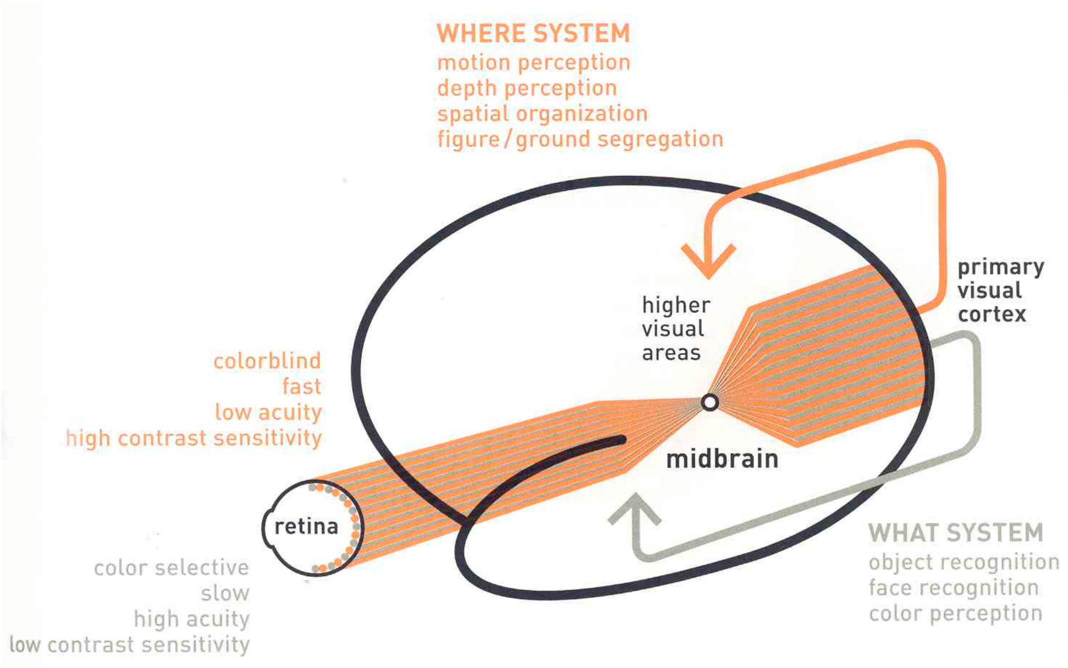

Diagram of the What and Where system
(Livingstone Margaret, 2002, 51)

A system does not need to be overcomplicated in its generation of objects. When recognising an image we hone in on the specifics. Given that perception was developed for survival it requires fast recognition by generalisation. Chess players of experience demonstrate this by being able to study a board and make decisions at a glance. They identify the key items to look at where as novices scan the board at random, taking much longer to assess the situation (Arbib Michael A., 1989). AARON already demonstrates this capacity by creating a simplified internal image of what it is creating. In The First Artificial Intelligence Coloring Book (Cohen Harold, 1984) Cohen describes the agency that governs the spacing of elements in the image as being cellular. It considers what is going on in a basic fashion. He also mentions in What is an Image? that he has modelled a shifting of attention in AARON based on what he believed the human art making process might be like (Cohen Harold 1979). Which I believe is corroborated by Arbib's statements.

An interesting feature of vision that led me to the consideration of neural nets is the feature of lateral inhibition. In vision this is known as center-surround and was discovered by Stephen Kuffler. The discovery was that small spots of light activated the ganglion cells in the eye better than large ones, being that cells inhibited their neighbours (Livingstone Margaret, 2002). This feature is demonstrated by Mach Bands (Arbib, Michael A., 1989), the bands of gray individually appear to have a gradient to them but are in fact an illusion caused by lateral inhibition. The geological layout used in the neural nets created by Teuvo Kohonen were what brought me to the consideration that a competitive perceptual model might be gleaned from such technology. Indeed Arbib says, "The first successful chapter in the theory of parallel computation for low level vision is probably that of lateral inhibition, the structuring of a network so that neurons inhibit all but their closest neighbours." (Arbib, Michael A., 1989, 118) This will be dealt with later.

Mach Bands
(http://www.nku.edu/~issues/illusions/MachBands.htm)

Given an understanding of some elements of biological perception we must discern what constitutes an image and how that would be verified as art. In Parallel to Perception Cohen states that art is not primarily beauty or decoration but the ability to convey meaning (Cohen Harold, 1973). Confirming and expanding on this Rollins states that a picture's standing depends on its causal history, that it also should be interpreted through diagnostic elements in it, perception is guided by these diagnostic elements (Rollins Mark, 2004).

By diagnostic elements we need not simply consider figures that may exist within the image. Although Cohen chose at the outset to give his program the ability to differentiate between figure and ground (Cohen Harold, 1979) Hofstadter points out that when we view a figure within an image we impose the ground ourselves. There are pictures that possess a distinguishable ground that can be isolated from the figures but there are also images that contain figures whose ground becomes unrecognisable when the figure is removed. What we are left with is a few chance lines that lose all meaning (Hofstadter Douglas, 1979).

Rollins also suggests a picture disassociates itself from words - this does not belittle the use of a programming language to investigate art. In a generative art application the use of an underlying language conveys meaning. Mapping of different structures on upon one another creates meaning. Such a successful translation of symbology would be referred to an isomorphism (Hofstadter Douglas, 1979). This allows us to view elements in an image as signs. The use of signs is essential for the translation of knowledge. It allows one to externalise experiences and grasp entities out of context, as recognising the sign of an object is not the same as recognising the object itself (Pagondiotis Kostas, 2004).

This justifies the abstraction of pictorial elements as algorithmic code that Cohen has used in AARON. But he is concerned that the machine in question not be a means of transforming a given image that he refers to as a “process” (Cohen Harold, 1974). The machine needs to create art on its own or meaning is lost,

otherwise the work is still a product of the artist and not the machine. Recycling the image will not invest any more meaning that did not already exist in the image. Seemingly contrary to this Goodman states that a copy (which is a process) is impossible, the representation diverges from the original. We accept the copy as a sign of the original not the original itself. He says that a representation is thus achieved, not copied (Goodman Nelson, 1984). But what Cohen clarifies on this is that like a photographic camera the computer is often a tool. We do not say that the camera took the photo, and thus the camera does not generate meaning, the artist generates meaning through his manipulation of the medium. A mechanical system capable of generating art would then require the ability to disassociate itself from the artist. We would say that the system is autonomous. Lacking autonomy the system will be nothing more than a tool. The machine that results from this investigation is something like the enchanted mop of the sorcerer's apprentice. It may fall short of a personality but it is free of its master and autonomous. One no longer perceives it as a tool but something that might be best described as a means.

Because one of the founding conditions of AARON was that it be denied any form of input Cohen has gone some length is disassociating the program in its infancy from anything that might be pinned down as a library of symbols. The machine generates forms from lines and these lines are created piecemeal. In this stage of the program's development it avoids a process definition.

In conclusion then we are not yet concerned with colour though that is not to say it is irrelevant, it is simply to say that we can omit it for the time being. The machine must be capable of abstraction of its own mark making. It needs to assess the creation of its image through quick and shifting attention. And it must possess autonomy to be considered an artist in its own right.

Chapter 2: Intelligence

At the heart of this paper is the need to discover pathways to forming artificial intelligences. The subject was first posed by Alan Turing who helped in the construction of Colossus during the Second World War. Turing proposed a game that could be played between a language using being and a human. The Imitation Game (Turing Alan, 1950) has become famously known as the Turing Test. The test amounts to the notion that if a machine can sufficiently fool a human being into thinking they are dealing with another human then said machine could be described as an artificial intelligence. The point of this argument lies in the fact that such a dialogue would eventually reveal the computer as a fake. The difficulty computers have with such simple concepts as analogies are illustrated by Douglas Hofstadter and Melanie Mitchell's Copycat program (Hofstadter Douglas, 1995).

The premise of Copycat is this: given the sequence of letters a, b, c changing to a, b, d, what would the sequence of letters i, j, k change to? As Hofstadter points out our immediate response is perhaps i, j, l. But a more smart Alec response would be i, j, d. But even that response would not be so obvious. How would a computer fair? Strictly bound to rules how can it be capable of the creativity that human beings are?

Intelligence in a machine is a contentious issue owing to the mathematics that the common computer is founded upon. Today’s fastest personal computers have developed far from when in 1938 Korad Zuse constructed a large machine called the Z1, from electromagnetic relays in his parent's living room. But they still rely on the binary math invented by George Boole. Tier upon tier of yes and no decisions create how common computers operate and are very different to they electrical switches in our own brains which operate on hundreds or thousands of connections between individual cells. Biology it would seem does not operate in binary.

Although it seems there is little hope of creating a form of artificial intelligence that is realistic enough to pass the Turing test, at this point I would like to propose the possibility of not a dialogue but a rhetoric. If, as in the case of an artificially intelligent painter one does not require a dialogue, can we infer intelligence if we can be fooled into believing the rhetoric of the automaton was the product of a human being alone?

One facet of intelligence is self-awareness. This is a central concern of the Parallel to Perception paper. A computer does not necessarily need to be entirely conscious of its elements, I am ignorant of the exact locations of my internal organs except when they cause pain and discomfort. In fact bodily function where a machine is concerned is not an aspect of intelligence (Hunt Daniel, 1985).

Systems of feedback are necessary for a machine to present qualities that may pass for intelligence. Cohen establishes this at the outset of Parallel to Perception. He also includes that adaptivity be an important aspect of the machine. Although AARON is a system which does not learn from image to image, the autonomous agents within it can be said to adapt whilst in the generation of images. In How To Make A Drawing he proposes a system that will have a form of long term adaptability. Unaware of this at the beginning of my researches my hope was to create a system that would do just that. Still, the current model of AARON does not seem to exhibit a behaviour of adaptability between images except in its evolution at the hands of Cohen.

I feel it is the ignorance of the real world that creates the problem of long term adaptivity for a system like AARON, least of all the limitations of physical memory which in computers - although we have superior storage in this age to when AARON was first proposed - is still inferior when considering a system that can adapt and learn over a period of years. I am assuming AARON still only adapts within the confines of generating one image. This supports the proposal of a non-process orientated image creator. But if we are to break Cohen's rule of keeping the input slot on AARON firmly closed it will have to be in such a way that the input is not merely

processed. If my proposal is nothing more than a glorified photocopier, which most computer processing methods are, then it will not possess the appearance of intelligence.

Cohen also rules out iterative digital art as intelligent. He cites Game Of Life as an example of a self-aware system that appears to exhibit intelligent behaviour. In the Game of Life system each living cell on a two-dimensional playing field monitors its neighbours and kills itself or generates more depending on the status of the surrounding cells as alive or dead. This system is certainly self aware but it is still a slave to its iterative elements. Hofstadter illustrates a very valid element of intelligence, that of boredom, or knowing when to quit a task (Hofstadter Douglas, 1979). AARON exhibits a simulacrum of this in knowing when to give up on an image. It lacks the full spectrum of intelligence of course that includes anticipation but that would require a long-term memory, which it does not yet possess. Game of Life exhibits neither anticipation nor boredom. It does not deviate from its rules despite self-awareness. AARON's self-awareness stems mainly from the aspect of what might be called chaos management. Like a real artist AARON is coping with an input of contradictory data that is fed to it via the computer's random number generator. This necessitates an adaptability and awareness that simulates intelligence. Were Game of Life to exhibit an element of negotiating chaos in its design then its self-awareness perhaps might tend towards the appearance of intelligence, rather than simply perpetuating the same iteration over and over.

That is not to say that iteration cannot be irrational or random though. In Stephen Wolfram's A New Kind of Science he reveals that simple automata are indeed capable of irrational behaviour (Wolfram Stephen, 2002). At the start of his book he demonstrates using some one-dimensional automata observing eight simple rules the capacity for complex behaviour. Of the 256 rules available to this system, rule 30 is the most striking as it exhibits random chaotic behaviour. Both AARON and Hofstadter's Copycat use a random number generator in their simulation of

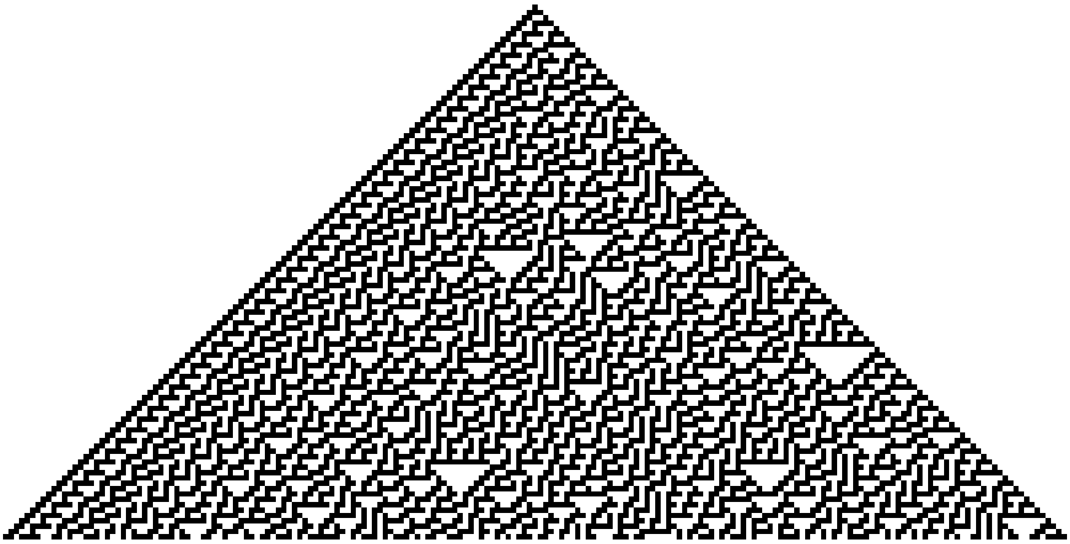

Stephen Wolfram's cellular automata exhibiting rule 30

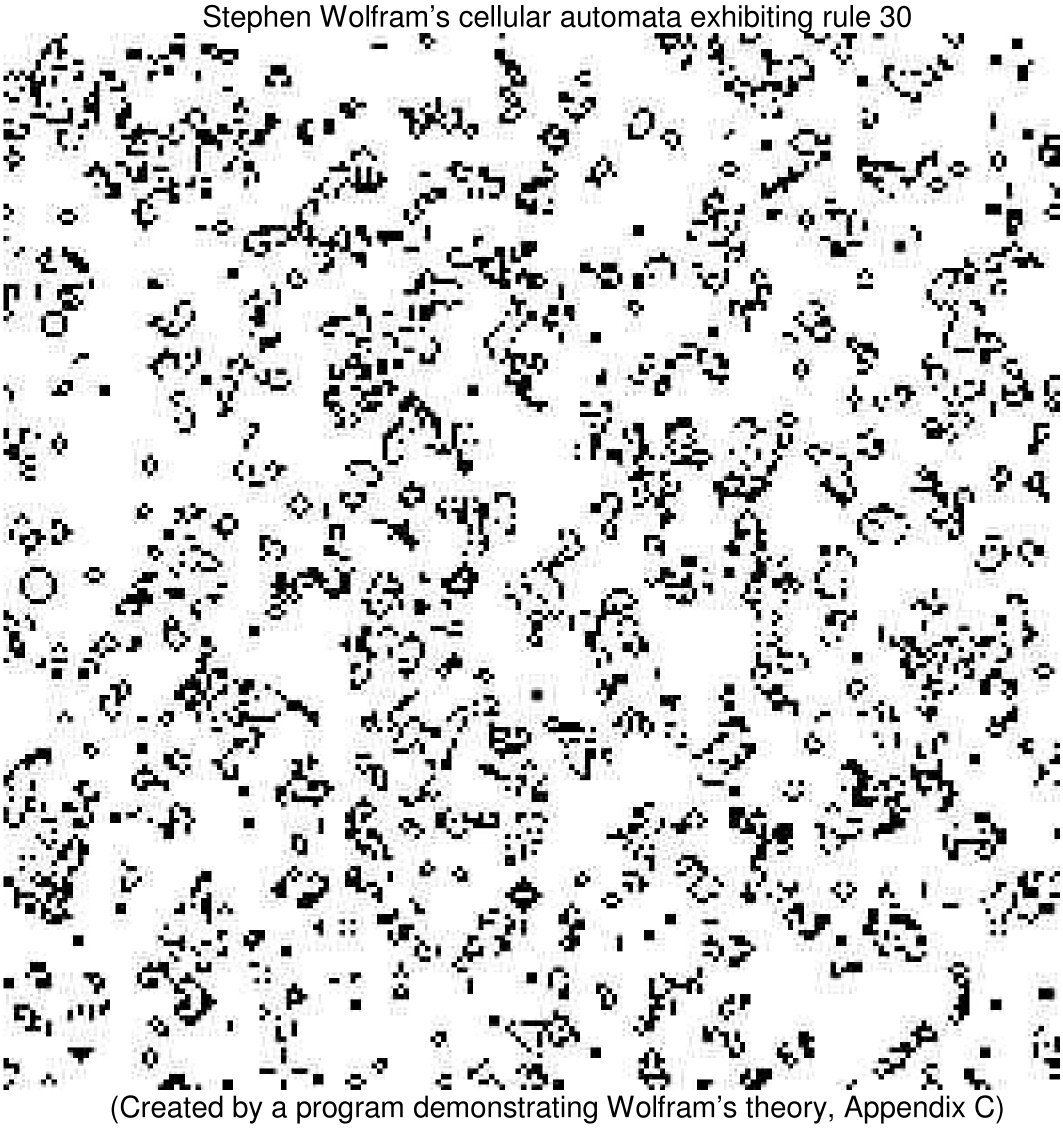

(Created by a program demonstrating Wolfram's theory, Appendix C)

An iteration of Game of Life
(http://processing.org/learning/examples/cellularautomata1.html)

intelligence and creativity. Hofstadter argues for fluidity, such as in water. An emergent quality that simulation of requires an underlying randomness (Hofstadter Douglas, 1995, 233). Wolfram points out that through the most basic logical conditions we can witness automata modelling fluid flow - and these automata exhibit randomness. Assuming a contradiction in theory I wrote to Hofstadter about this (Appendix B) and he replied that fluidity can be emergent in many systems where there is an underlying irrationality, be they based on automata or a random number generator. The value of the automata seems that we are not limited to just one type of randomness.

In developing an artificial intelligence we must base it on heuristic search and domain knowledge. Heuristic search means the search through a pyramid like tree of decisions ending in goals or terminal nodes (Hunt Daniel, 1985). This structure is not unlike many of the neurons that exist in the brain. AARON already exhibits this quality of expert knowledge and heuristics.

AARON uses a language that is implicitly heuristic: Lisp. Lisp stands for List Processing and was invented by John McCarthy in the years 1956-58. Hofstadter describes it as a prime choice for the pursuit of artificial intelligence (Hofstadter Douglas, 1985) and details its operation. The language is broken into two primary components: atoms and lists. An atom is like a reference device and the list is like a structure device. An atom may refer to values or lists and a list may contain atoms or lists. This almost confusing aspect of Lisp gives it a tremendous power to imitate the structure of human associative thought. Because one can create recursive heuristic pathways of logic in Lisp it offers clear advantages over other languages in modelling human intelligence.

In What is an Image Cohen details some of the paths of logic that AARON takes in making a drawing. The tree is to some degree recursive, feeding back on itself till satisfying requirements.

In 1956 a meeting was held amongst scientists researching artificial intelligence and their hopes were particularly over optimistic in their certainty that we should be by now having the majority of our tasks being performed by automatons. It is with a great degree of modesty then that I embark on researching solutions to such problems and I realise the most profound difficulty will be satisfying the necessity that the outcome of my researches will be more than just a tool. The key factor in AARON's intelligence I believe is its pursuit of chaos management. The intelligence we see in other life forms comes from the fact that they too exhibit a readiness to cope with all manner of calamity. If I propose the opening of AARON's input slot, then what must go into that slot must be equivalent to the outpourings of a computer's random number generator. Or in light of Wolfram's research we might conclude that perhaps there are simple rules that will effect an irrationality we can use. Our input must not be processed though, I could only suggest that somehow it would be meditated upon. And above all, the output of such a device need fool the layperson that the produce is of a human hand if we are to effect an intelligent rhetoric.

Aside from intelligence is the matter of creativity. In How to Make a Drawing Cohen offers the supposition that creativity is realignment of ones internal representations. Similar to this in a recent article in the New Scientist by Helen Phillips (29 Oct 2005) we find that noradrenalin levels in the brain subject an effect on creativity. Noradrenalin hinders the ability for neurons to 'talk' to each other in the brain. Giving someone a precursor of this chemical seems to hinder their ability to solve creative word puzzles. This would seem to confirm Cohen's assertion in a way, but to clarify we would say that a high level of intercommunication would be desirable if our machine is to present itself as creative. There is not a Turing test for creativity and in his last paper available on the internet Cohen says that he now avoids use of the word creativity or the attempt of defining it (Cohen Harold, 1999). But in light of the scientific evidence we can say there is certainly an analogue of creativity in

AARON for its intercommunicative aspects, and if we should model a creative machine then we should take after this.

Chapter 3: Digital Art

The advent of digital technology allows us to reassess what artists create now that a platform is available for creating work that can - in terms of computer code - be replicated perfectly. AARON's principle advantage over its creator was its tirelessness, capacity for omnipresence and its ability to create a quality of line that as Cohen put it seemed better than his own. As he moves from AARON's initial incarnation towards a model that created representational images Cohen remarks that he has given up drawing because AARON seems to do the task so much better than him (Cohen Harold 1982). Before considering the technology I have researched to apply to AARON I will review a number of digital artworks that operate in a parallel manner to Cohen's work and consider the ramifications of working in a digital medium.

The most immediate parallel in terms of the automaton that exists in the gallery space making works of art is that of Chris Burden's When Robots Rule: The Two-Minute Airplane Factory. Using a design based on a model kit produced by Peck-Polymers Burden created a fully automated model airplane factory. The machine, like AARON, requires little more manpower than is required to load the raw materials into the factory. The machine then does the rest, even trumping AARON in one respect by winding the elastic powered propeller of the finished plane and launching it to spiral up into the space of the Tate Gallery where this exhibition took place. It has no intelligence of course but like AARON the airplane factory makes an allusion to the assembly line, and industrialisation. Burden calculates a total of 71,000 man hours saved in the process of the whole exhibit considering the machine can build and launch a plane in two minutes, whereas Burden admits that he himself might take up to two-hours to construct one of these model airplanes. AARON was certainly not conceived as a time saving device but through out his papers there is the sub-script mention of the speed and tirelessness of the machine. Cohen is saved

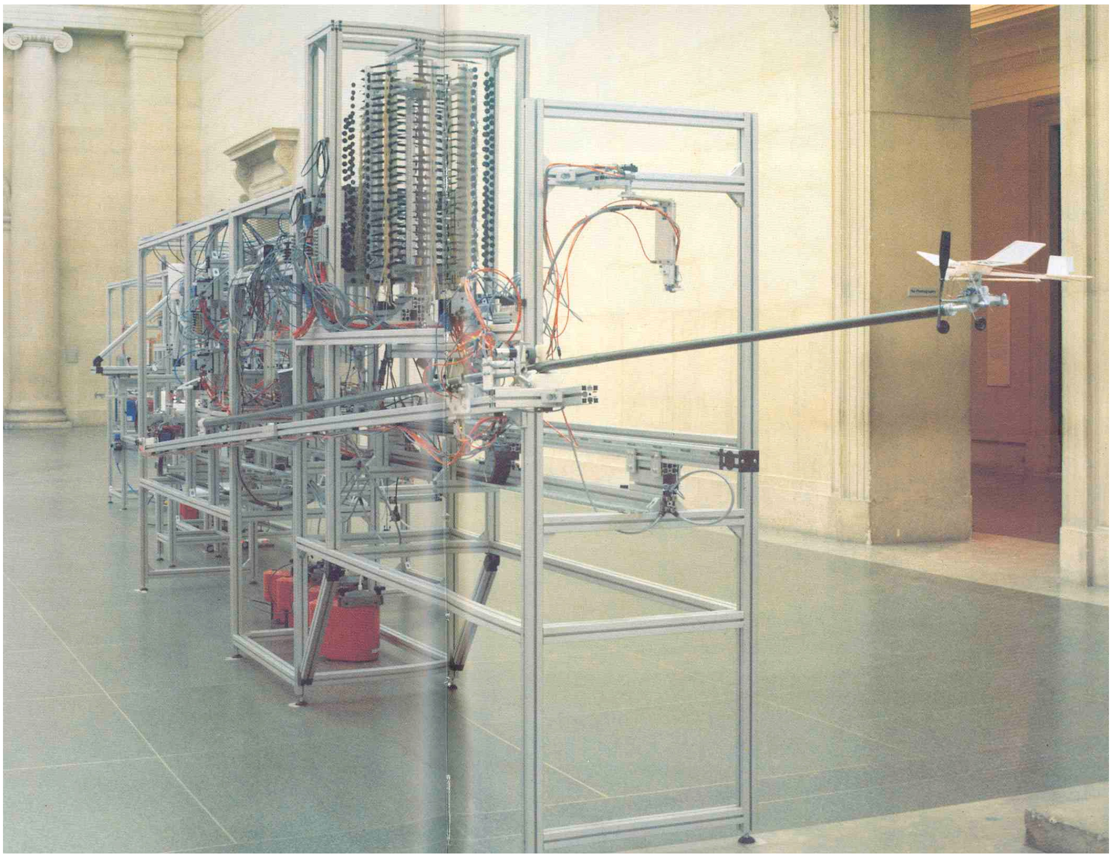

When Robots Rule: The Two-Minute Airplane Factory 1999 Chris Burden
(Morris Frances, 1999, 51)

many man hours and can effectively operate as an artist in many places at once, just as Burden's factory tirelessly launches airplanes which are retrieved and sold as part of the exhibit (Frances Morris, 1999).

A reversal of the process of AARON is Keith Tyson's art machine. Whereas AARON is infused with Cohen's understanding of art and left to create, the Artmachine of Tyson's operates as the creative impulse. AARON is an exercise in autonomy and the Artmachine exists to deprive Tyson of autonomy. They are strangely two sides of the same coin because both use the random number generator of the computer as the feed for ideas. The random number generator in AARON serves as the friction of living but in the Artmachine the generator serves as a sort of I-Ching. Tyson gives recourse to his mechanical oracle to divine what he may create. The output of this machine, although not detailed in its operation is quite bizarre: a sculpture of a hammerhead flamingo, a drum kit for accompanying the Dawn Chorus, the casting of a Kentucky Fried Chicken Menu is bronze and even the stalking of a client who purchases a work named "The Entourage" who receives the produce of the sudden attention he is warranted by paparazzi, a documentary film crew, a private investigator, etc (Tyson Keith, 2002). Not all Artmachine iterations go to completion but even the suggestion of the work becomes a work of art in its own right. Tyson's Artmachine adds another possibility to expanding on a system parallel to AARON, that of the employ of assistants. Although Cohen describes AARON as not a tool but an assistant (Cohen Harold, 1982) he is in turn an assistant to AARON, having to feed it ink and paper, much like the gallery assistants feeding Burden's airplane factory materials to make planes from. Art that questions autonomy is intriguing in the fact that the piece gains the opportunity to make use of people. This perhaps is another element we require, the machine becoming the master.

The mastery of the machine was suggested in the film Alphaville of 1965 where people are forced to live sterile lives subject to the whim of a machine. Angela Bulloch's piece Betaville of 1994 takes its name from this film. It is a grid of horizontal

Iteration :	AMCHII• XXXIX	**Size / Duration :**	apprx.150 x 180 x 75 cm
Title :	Hammerhead Flamingo	**Untitled number :**	•
		Series & Editions :	Unique
		Hanging Specs. :	Floor
		Location / Site :	•
Format :	Sculpture	**Documentation :**	•
		Date :	1995
		Conditions :	Lit from below
Status :	Proposed	**Other variables :**	•

Brief description of proposed work / Reproduction of finished work

AMCHII•XXXIX
"Hammerhead Flamingo"

Notes : This rubber and plastic sculpture stands on a circular lightbox as the piece needs to be lit from below, therefore, the piece should be exhibited in a dark or dimly lit space.

Artmachine Iteration: Hammerhead Flamingo 1995 Keith Tyson

Iteration :	AMCHII • XX	**Size / Duration :**	Various
Title :	The KFC Notebooks and the UCT. (Untitled Conspiracy Theory)	**Untitled number :**	•
		Series & Editions :	Unique (while sustaining 3 edi.)
		Hanging Specs. :	On platform against a wall.
		Location / Site :	Any
Format :	Mixed Media Installation	**Documentation :**	Prints, notes, computer disks,
		Date :	1995
		Conditions :	•
Status :	Proposed, fabricated & exhibited	**Other variables :**	Platform to follow form of wall

Brief description of proposed work / Reproduction of finished work

(Installation Details – Anthony Reynolds Gallery 1995)

Notes : This piece is a small installation that sustains a series of multiples and prints.
The work consists of a platform, computer terminal running a viral animation, notepapers, a soundpiece, a set of 12 prints, several correspondences, two bubble jet prints (one encased in leaded glass) and the complete Kentucky Fried Chicken Menu cast in lead (ie. all the burgers, dippers, fries, family buckets etc.)
The whole piece is based around a conspiracy theory with Colonel Sanders and the KFC at its centre.

(Tyson Keith, 2002, 55)
Artmachine Iteration: The KFC Notebooks 1995 Keith Tyson
(Tyson Keith, 2002, 53)

and vertical lines created by a robotic armature sensitive to a switch in a viewing seat. It was meant to evoke the feel of a city from a plan view, reflecting the city of the work's namesake. The lines turn vertical when a viewer sits on the bench before the work and horizontal otherwise (Bulloch Angela, 1998). Bulloch remarks that it makes the process of viewing visible which Cohen has also done in creating a turtle to draw large-scale images made by AARON. He even adapted the drawing routines to suit the nature of the twin stepper motors which would veer the pen to its destination in the same way the human hand veers the pen to a location, with the same lack of exactitude (Cohen Harold, 1979).

"[With Betaville] I'm making another kind of picture than a traditionally recognised painting... but the way it works does question the idea of the artist as being the creator of something original." (Bulloch Angela, 1998, 98) This perhaps is an aspect of most digital art these days given the advent of the internet where Java and Flash have given the end user as well as games and applications the opportunity to explore different forms of digital art. Maeda reiterates in different ways the dynamism given to us by digital media, from the different colours and complexities one can explore from the initial concept to interactivity and motion (Maeda John, 2000). Even the surface is dynamic he tells us, as we are using digital tools with screens that imply an unrolling terrain to explore. And although it is devoid of interactivity Cohen remarks that AARON produces endless variations, each original. Bulloch's work seems to contradict this in its user participation, but it has the same aspect of chaos as AARON, we cannot predict the behaviour of the visitors of her work and thus the final image. I think that opportunity here is not for originality as such but to blur the distinction as to whom the artist is. Cohen suggests the artist makes the art to redefine himself and art (Cohen Harold, 1974). What better way to do this than to ask whether the artist is at all part of the picture?

It is interesting to note Cohen's remarks of comparisons made between AARON's work and his own (Cohen Harold, 1979) and that he never intended for his

Betaville 1994 Angela Bulloch
(Bulloch Angela, 1998, 98)

work to replicated by AARON. Certainly not in the fashion of A. Michael Noll referred to by Maeda (Maeda John, 2000) and Hofstadter (Hofstadter Douglas, 1985). Noll created simulacra of Mondrian's work digitally by analysing the distribution of the forms in the piece Composition With Lines. My own intuitive eye selected the real Mondrian from the impostor but when the work was shown to a group of people many were fooled into thinking the fake was the real one (http://www.citi.columbia.edu/amnoll/CompArtExamples.html). One can see something of the same in looking at the paintings of Cohen's retrospective of 1965 at the Whitechapel Art Gallery in London. John Richardson refers to Cohen as having a preoccupation with duality and ambiguity (Whitechapel Art Gallery, 1965) - Cohen himself says he sought to impose neutrality upon the canvas.

Although unintentional we see the same neutrality being meted out upon the canvas (AARON fills its workspace) and the duality and ambiguity is now played out in the sense that who the artist is now ambiguous. One could also suggest he has reached further heights of duality as well in becoming two artists as opposed to one.

This is something no doubt I will be unable to avoid in a creation of my own. I may succeed in creating an independent hand to draw with, but it will still be noticeably mine.

Computer Composition with Lines 1964 A. Michael Noll
(http://www.citi.columbia.edu/amnoll/CompArtExamples.html)

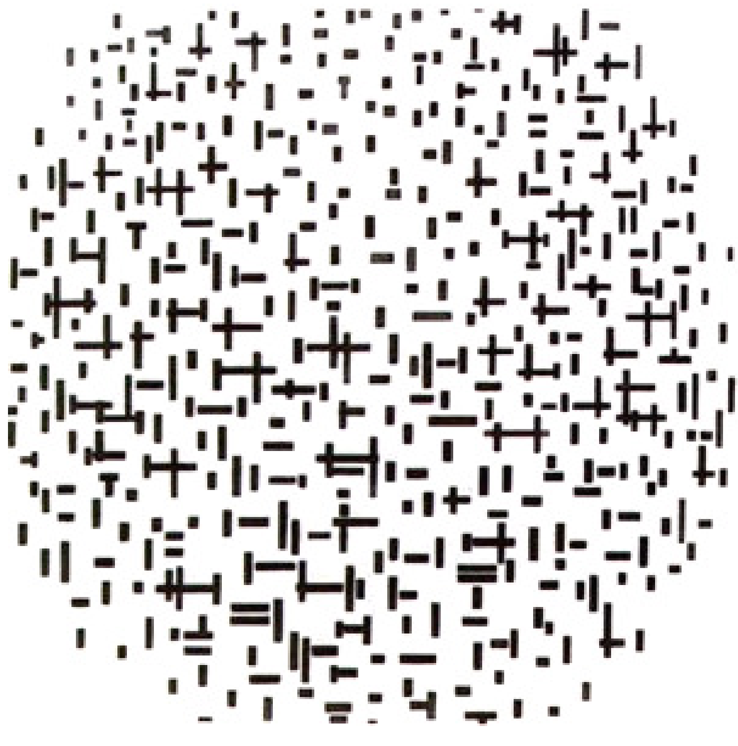

Composition with Lines 1917 Mondrian
(http://www.fiu.edu/~andiaa/cg2/chronos.html)

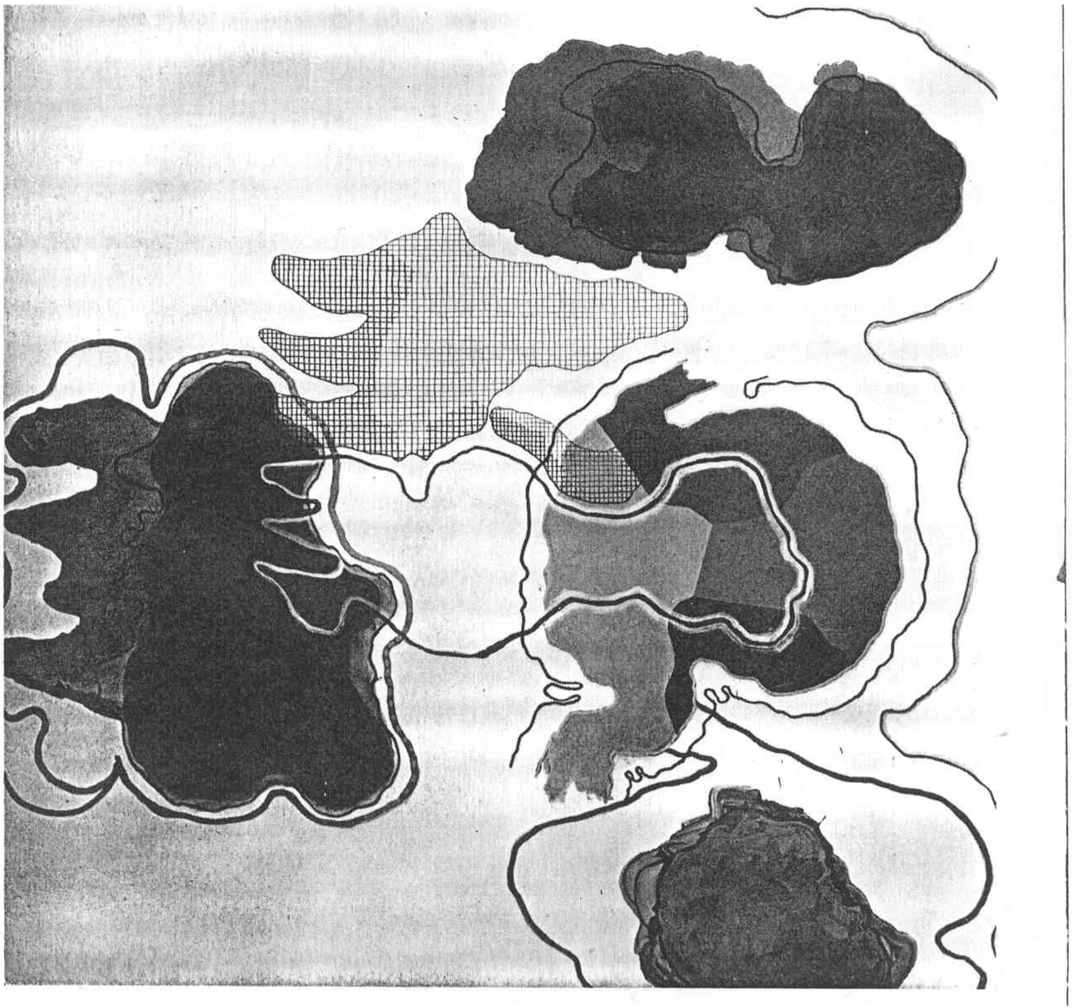

Eunice, Harold Cohen, 1964

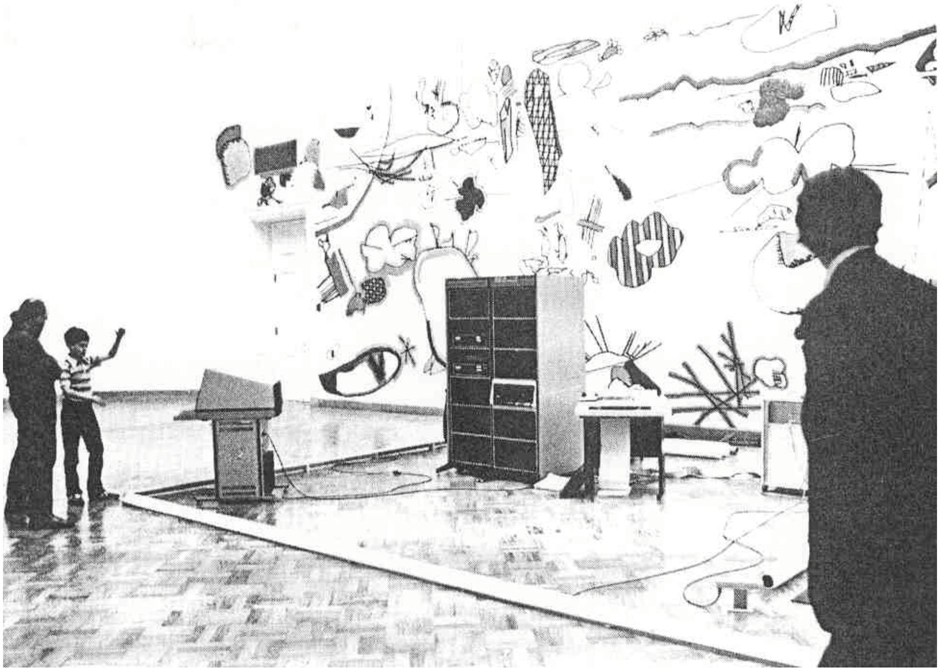

(Whitechapel Art Gallery, 1965, cat. 55)
Drawing, Harold Cohen, 1979
(Tate Gallery, 1983, 27)

Chapter 4: Neural Nets

I have researched neural nets for creating an artificial artist because they are a biological analogue. I will review the various neural net technologies I have come across and whether they are appropriate for this use. In contacting Harold Cohen he responded that he had no intention of using neural nets as they were a means of determining what is best and the pursuit of art cannot be reduced to such a thing (Appendix A). I agree that if one begins to say that art must be *this and that* one has lost sight of what making art is. The problem here is to see if neural nets can be used without reducing them to achieving a specific goal. The purpose is to see if their adaptiveness and their ability to generalise can be put to service.

Before considering the choices it pays to know exactly what a neural net is.

The brains of animals are composed of many nerve cells called neurons. The neuron receives stimulus from a branching network of *dendrites.* Stimuli travel down the dendrites to the *soma*, the body of the cell. Given enough stimuli the neuron fires a charge along another process called the *axon*. This branches out to connect to the dendrites of hundreds of other cells.

Neural nets are based on this model using a collection of threshold dependent switches. Numerous channels of data enter the artificial neuron, if the data exceeds a certain value the neuron 'fires'. Reflecting the weakening and strengthening in connections between biological neurons the input channels to the neuron exert a certain magnification or reduction on their data. These are referred to as the neuron's input weights and are central to neural net theory. Over a period of time that for many neural nets is considered their training period, they learn how to respond to various sample data. Once the net is trained and the weights have been adapted to the desired output from the neurons it is ready to classify data.

I initially experimented with Kohnonen Self-Organising-Maps (SOM). These types of neural nets invented by Teuvo Kohonen are based on the principal that

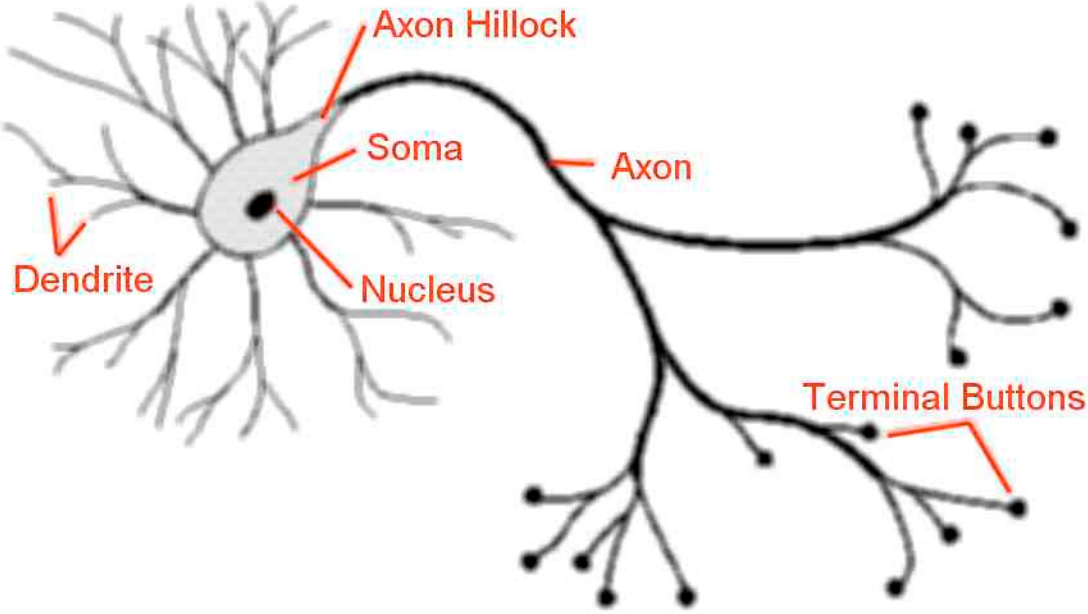

Schematic of a biological neuron
(http://vv.carleton.ca/~neil/neural/neuron-a.html)

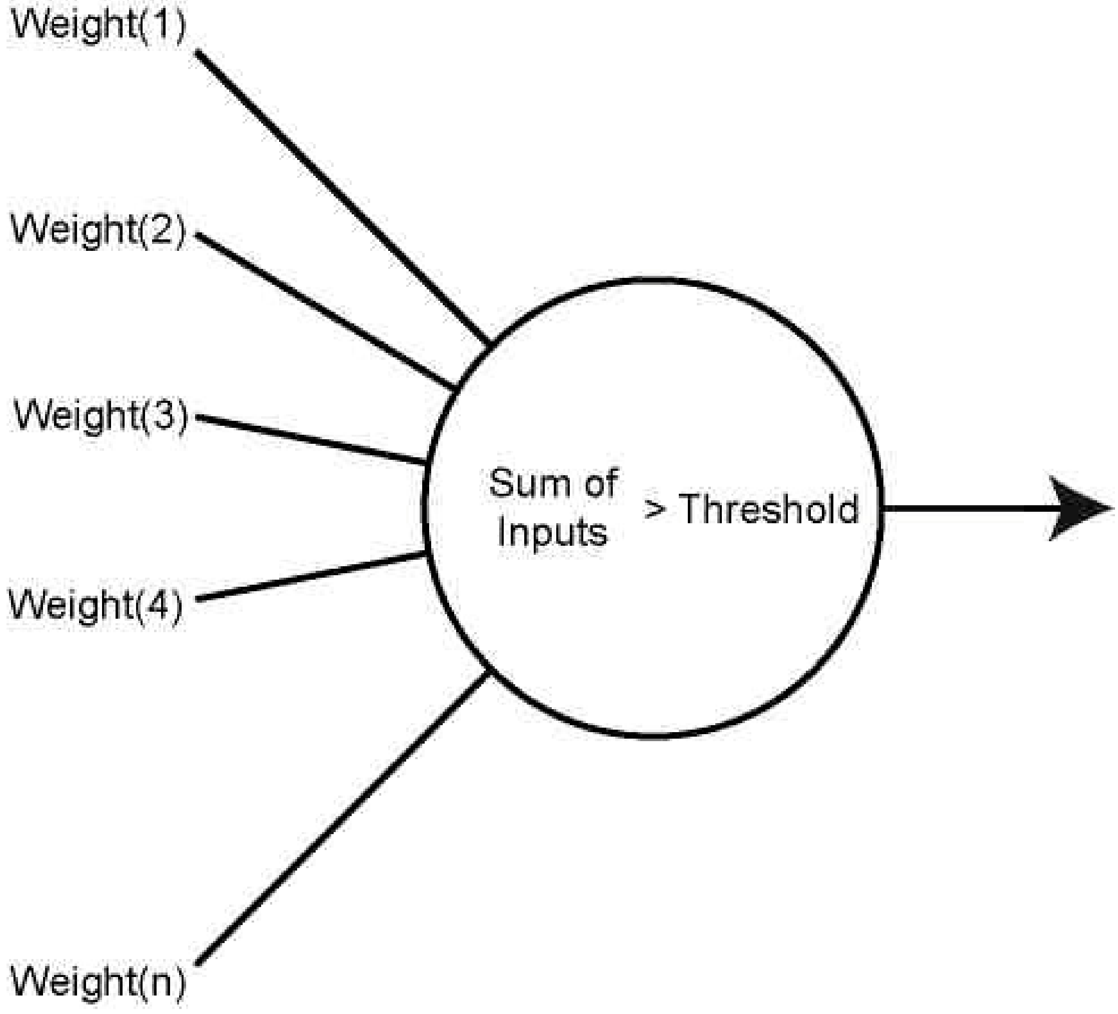

Schematic of an artificial neuron
(Based on http://www.iiit.net/~vikram/nn_intro.html)

distance between neurons in the brain plays an important role, especially the physical location of these cells. I was lead to investigating the use of Kohonen's algorithms because of the competitive training method they use. Having learned about lateral inhibition I believed at the time that a neural net that exhibited this sort of behaviour would be of benefit. It would later become apparent that I had completely misinterpreted competitive learning for lateral inhibition - the two being at odds.

The first of these algorithms that I explored effected the clustering of similar data together. The use of this can be applied to gathering data with similar features together. When this type of net is trained upon an image the net gathers colours of likeness together. When one wishes to begin learning how to program neural nets this is without a doubt the easiest algorithm to learn and it is recommendable that regardless of one's long term goals for programming neural nets that they start with the clustering algorithm. It is what we programmers would call a, "Hello World", program. When one begins learning a new language one must first elicit the most basic of responses from the language. This has the effect of breaking the ice with the computer in said new language - the importance of which being that you will not get very far if the computer cannot give you a response. The, "Hello World", program usually takes the following format:

```
print "Hello World"
```

The clustering algorithm is by far more complicated in comparison but it is the most basic of responses one can elicit from more than one neuron. Simpler programs utilising neural net technology are generally concerned with the operation of one neuron only, usually demonstrating a logic gate. A single neuron cannot be considered to networked with no other neurons to pass data to. The hardest thing comprehend in developing a neural net is how the neurons are interconnected. I was

fortunate enough to find an online tutorial during my research that explained piecemeal how to construct a clustering algorithm (http://www.ai-junkie.com/ann/som/som1.html).

Describing the algorithm in the most basic terms; it consists of a grid of neurons. Each neuron has the same number of inputs and weights. These weights are set to a random small value before the training begins. Each training cycle a random selection is made from the sample data and the neuron best matching that data has its weights altered to resemble that data. Not only does the selected neuron get altered but also a radius of effect is cast around that neuron. All neurons on the grid within a given Euclidean distance are altered to be like the neuron at the center of the circle of effect. The training cycle begins anew and the next time round the circle of influence gets gradually smaller.

Say our data is a selection of colours. As the neural net trains itself on these colours what happens is that similar colours will find themselves grouped together as the net continues to pick the best matching neurons.

As I had the algorithm to hand and I had put a basic shape drawing routine together I combined the two to see what would happen. Try as I might to see what purpose there might be to it, I found none. The clustering algorithm generates zones of data similarity at random. It is impossible to predict where or how it will demarcate these zones. Although there is a concern for composition by the net it is not consistent.

After the clustering algorithm I tried the geological SOM. This is a kind of elastic medium that learns to adapt to a series of points that are given to it. These may be randomly selected from within a given area or they may be a specific quantity of points. I must credit the original author of the algorithm I adapted for experimentation. Lacking a down to earth guide for making this particular kind of neural net I adapted code I found on the internet demonstrating the SOM (http://www.patol.com/java/fill/index.html). Adapting a series of formulae into a

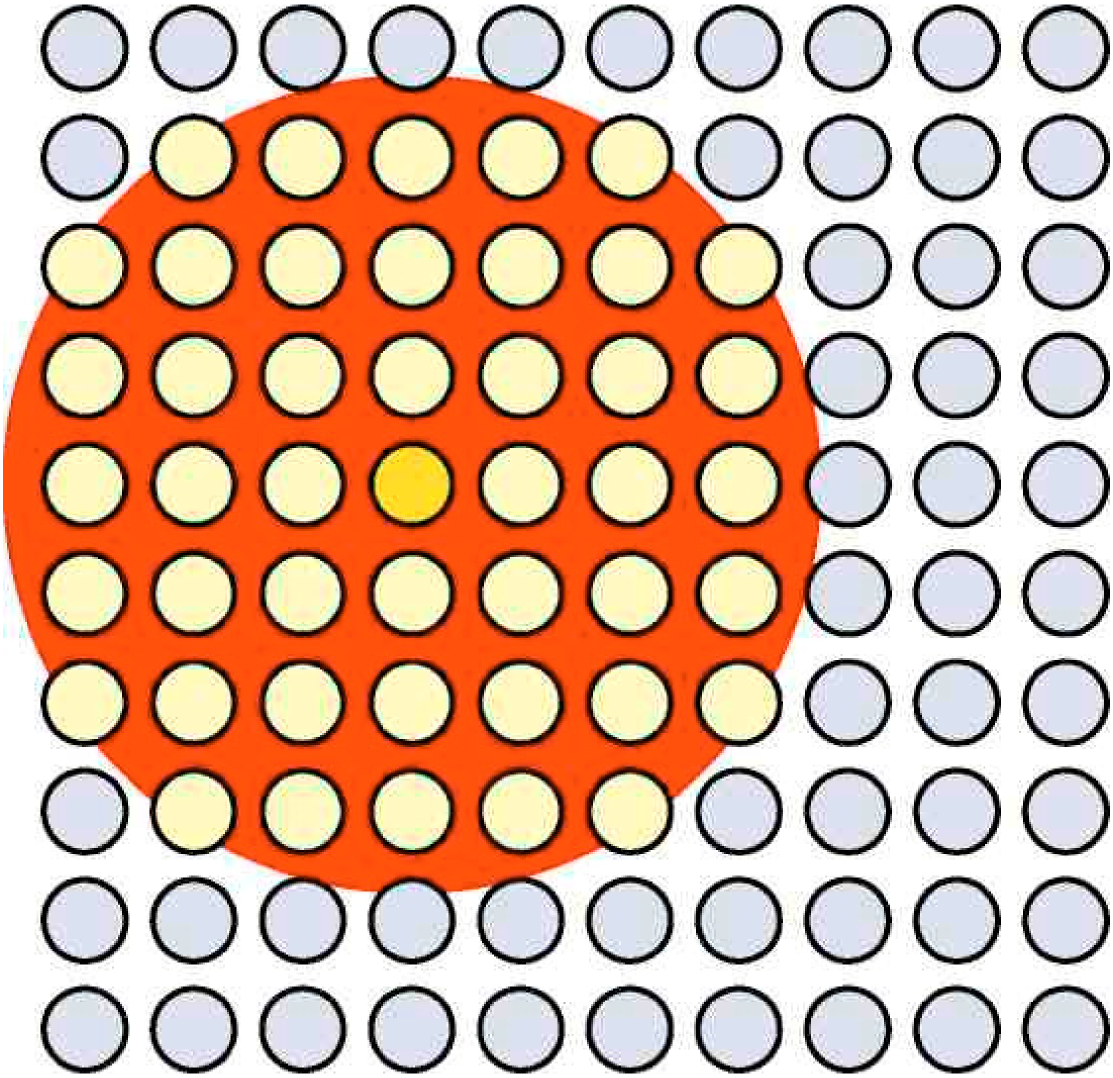

Diagram of how the neurons in a Kohonen clustering algorithm influence their neighbouring cells

specific language is generally quite a long and drawn out process. Luckily for me this particular code was written in a language familiar to me. The specifics of the training principles were unclear but knowing the semantics of good programming I was able to strip out a lot of unnecessary code and make the operation of this code much faster. I have made good efforts to ensure that the original author of the code has been credited in the different versions of this code I have published on the internet. This is far from a plagiaristic operation as the code at hand was offered from the given site for use in the public domain, and I in turn have done the same with not only this particular code but all of my work to date.

The net is yet again a grid of neurons but the net occupies an imaginary physical space. During the training we give the net one of the points we wish it to learn - a point in the physical space the net occupies. The net contracts around this point. The closest neuron to this point moves towards the point in space and pulls the rest of the net along with it. The points given to the net can be given to it randomly or cyclically. As the training continues the area of influence the best matching neuron exerts begins to fade (in a similar fashion to the clustering algorithm). Smaller areas of the net are dragged towards the given input. By this time the net has settled over an area that covers most of the points input to the net and individual areas of the net are stretching to best ensure the net covers all the points evenly.

When one uses a string of neurons as opposed to a grid this algorithm forms a solution to the Travelling Salesman problem. This is the problem of a salesman who must visit a number of cities but must take the shortest sequential route. Since the geological SOM tries to distribute itself evenly across a number of points this algorithm offers a solution (http://www.patol.com/java/TSP/index.html).

Having access to this algorithm I believed I might put it to use somehow. I thought perhaps it might offer variations of theme within an image. I set up a program to draw a random polygon using only points on a grid. I then trained the net on that polygon and the edges of the drawing space. Having achieved a distorted canvas

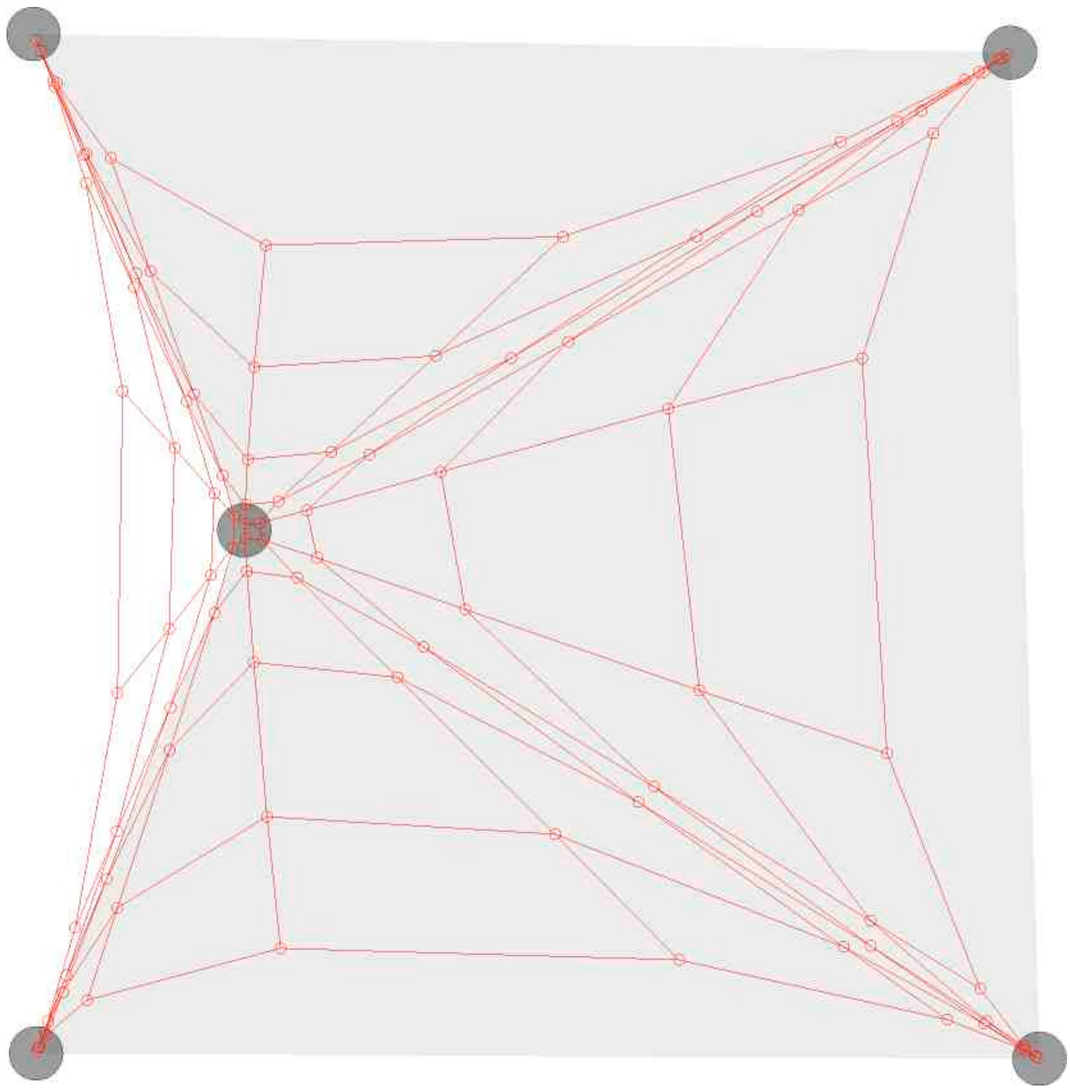
Example of a trained geological Kohonen Self-Organising map

grid from this The polygon would redraw itself using points from the new grid but to only allow its vertices to shift to points on the grid nearby. I was very unsure of the value of this model. I thought perhaps that if AARON were using this system it would draw ever similar images. What use that may be I cannot say and I had not thoroughly read Cohen's papers at this time nor was I aware of them. Had I been I would certainly conceived differently and it may why I received such a negative response from Cohen in contacting him and demonstrating this model.

Before adapting the geological SOM I had fortune to encounter another site demonstrating neural nets for the purposes of pattern recognition. Adapting code from this site (http://richardbowles.tripod.com/neural/hopfield/hopfield.htm) I was able to see the operation of a Hopfield neural net. The Hopfield net is notable for its ability for pattern reconstruction. Giving it a broken pixelated image one can execute the training of the net and within two cycles the net will have reorganised the data to resemble one of its stored images. The Kohonen networks are different in their training methods for they do not require any supervision. Most other networks require a selection of training samples to guide its behaviour. The images in its store need to be carefully crafted though and when I experimented with this net I found it unable to deal with images larger that eight by eight pixels. Further more it could only process them in black and white. Although I am limiting my investigation to black and white I would not want to take on a system that would never have the capacity to learn how to deal with colour. It is an area for continuing research for Cohen and I'm sure I will want to do the same in the future.

A further problem in pixel based systems is that a good pattern recognition system should be able to exhibit *translation invariance* (Bishop Christopher M., 1995). Let us say for example that I wanted to classify a pixelated image of the letter "a". My training sample will no doubt be placed centrally to the grid of pixels. If I am classifying something that is human input then I may well have a duplicate of the pixelated letter "a" but it will likely be off center. Although there is a statistical

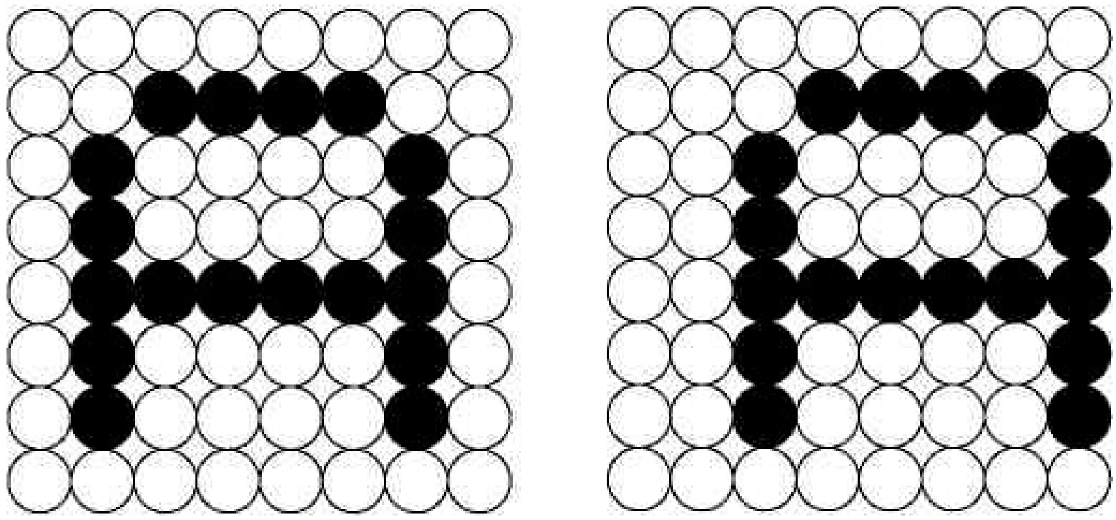

Example of what a translation invariant system of pattern recognition would need to overcome.

likelihood of recognition occurring, the computer - lacking the capacity for analogy - will report some error.

A way around this problem is not to use pixels but vector based representation instead. A vector representation will be dependent upon a series of points that will constitute a path. Because the vector system will be concerned with relative distances between the points it will make no distinction as to scale or positioning on the canvas being examined. A neural net responding to vector paths was what I worked on next after the Kohonen and Hopfield algorithms.

Although I had learned some fundamental concepts in trying to apply the Kohonen neural nets to this problem it became apparent that they were largely inappropriate for the task. The lateral inhibition I thought the competitive nature implied was in fact lateral confirmation. The net did not suppress signals around neurons best matching the given data but strengthened them. I have experimented with trying to invert this principle but the results are abstract and freakish. Their use is interesting in creating artistic contradictions perhaps but for the purpose of serving an electronic mind they are irrelevant.

I turned my attention to a form of network that is composed of three layers of neurons, transmitting their data from the input layer, to a hidden layer and then to an output layer. These types of network are known as a feed-forward networks as they send data in one direction through the network. The specific type of feed-forward network I used was a Back Propagation Network (BPN). The reason for this is that like the Hopfield network it depends on a selection of data to inform its training cycle. During the training period the network monitors the difference between the desired output from the network and its actual output. This re-feeding the network the information gives it its name. Once the error in the output is appreciably small the net is ready to classify data.

Using an open source BPN available on the internet (http://www.patol.com/java/NN/index.html) and guidance from a description of how

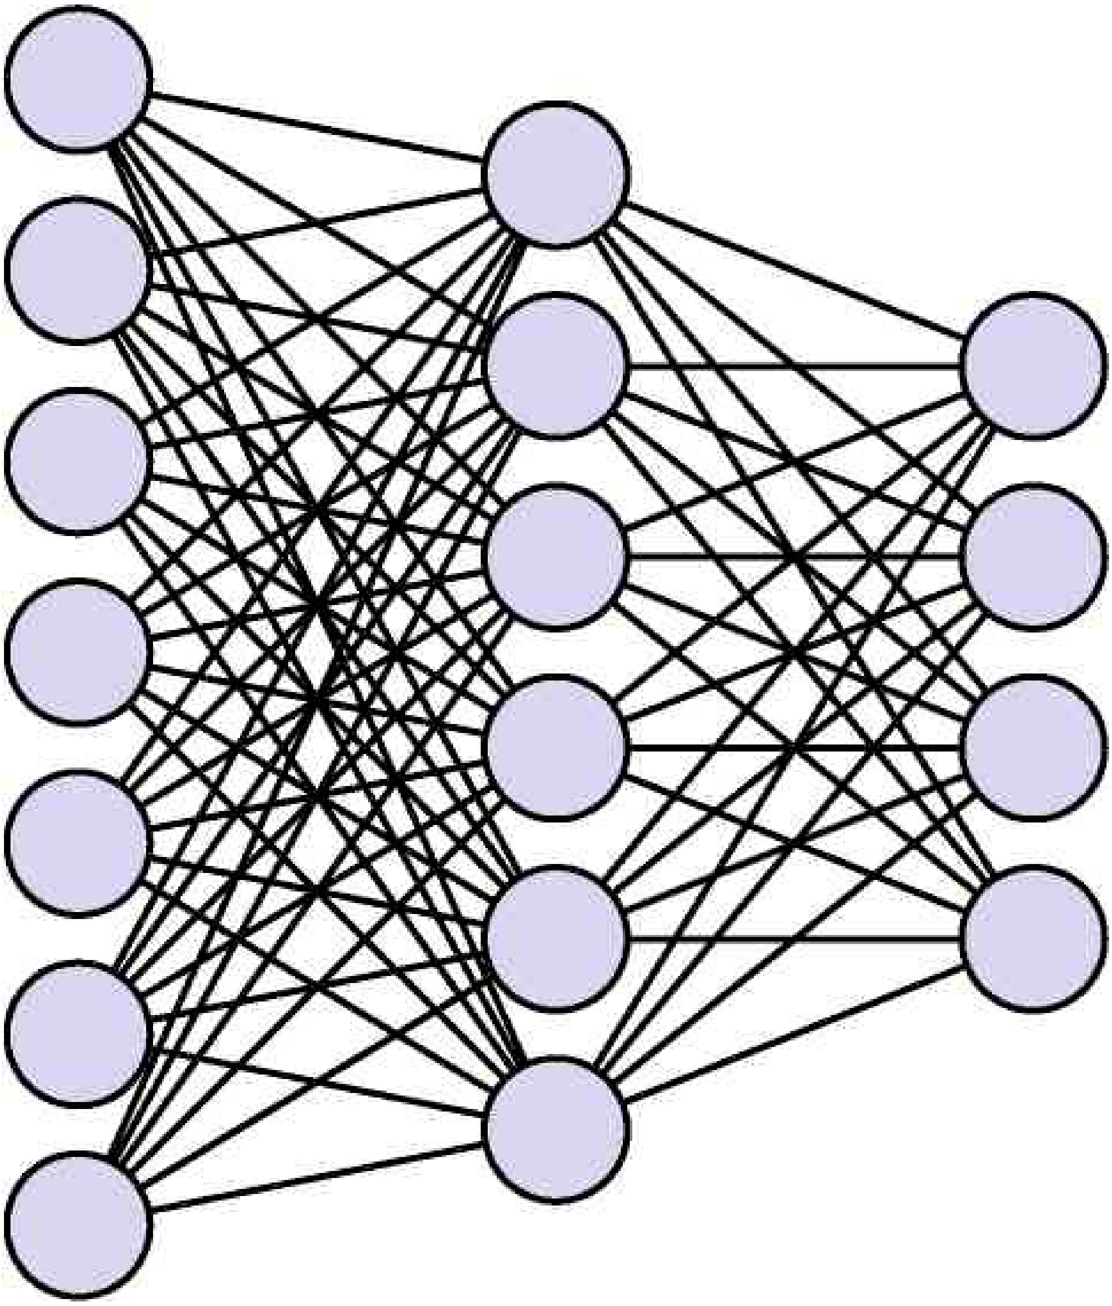

Abstract diagram of the structure of a back-propagation neural network

someone made a gesture analysis system (http://www.generation5.org/content/2004/aiTabletOcr.asp), I constructed my own gesture analysis program.

The motive behind this was that perhaps a system parallel to AARON could learn by indoctrination. Cohen teaches AARON by using collections of algorithms. Effectively he has to unmake and re-make AARON in order to educate it. There is already mention of an editor for inducting new shapes into the system in the paper How to Make a Drawing (Cohen Harold, 1982). Whether this is still in use or made its way fully off of the drawing board in AARON's current incarnation is unclear. The gesture recognition system I would create for an artificial artist would be used to learn strokes given to it by a tutor - this could be anyone. The neural net would then classify what kind of stroke it is, whether it is similar to its current store of strokes and marks or whether it has a case for classifying a new mark. The program's understanding of mark making would then slowly expand as it is taught new marks.

The gesture recognition model I have built is very basic and takes input to the network in the form of a series of points on a path drawn by the user. The program then reduces the number of points to a specific quantity and calculates the angle between each point. These angles are fed to the net's input layer. The output layer consists of one output per gesture. The implication of this is that the network would have to be rebuilt and retrained whenever it identifies a new class of mark. One cannot simply tack a new output neuron on to the network owing to the fact that there is total interconnection between one layer and the next. The pathways the net has meted out will have to be rebuilt.

The problem of using the gesture recognition system as a means of tutoring the system raises the problem of it becoming a process application. Is the system being taught or does it just accrue a bank of simplified images that it combines - becoming a means of collage rather than intuitively feeling out an image as AARON does? By having a store of strokes it would operate in a different fashion to AARON

who in its original form created marks without any preconception as to what it created.

But AARON's later incarnations offer some hope for the use of such a method as it is concerned with the generation of figurative images containing people and plants. This suggests a level of preconception and perhaps we can take the tutored marks and strokes as guide for the system's behaviour. Such a guide would have to be fairly loose in order to maintain a freehand character to the marks. Although imperfections would still exist in the tutor’s hand the repetition of marks would eventually evidence a lack of intelligence on the machine's part.

The Kohonen clustering algorithm might even be of some use for a mark-making tutor by being able to group marks of similar qualities together.

My practical research in this area has come up short but as my research comes to closure I have found uses of neural nets in the fields of creativity. Dr. Stephen Thaler has built a neural net system that resembles a combination of two feed-forward networks. His Creativity Machine is trained upon data and then it is internally “tickled”. What this means is that some of the connections in the network are randomised to create new ideas (http://www.imagination-engines.com/cm.htm).

Ed Burton who has created an online playground of virtual spring based Meccano is also using neural nets to create an adaptive system for representational drawing called EOR (Emergence of Representation). “EOR will have no prior knowledge of drawing; it will scribble with its virtual arm and observe the resulting marks with its simulated eye. Primitive drawing skill will develop over time through the dynamic self-organisation of a unique architecture of neural networks linking both arm and eye in a feedback loop of spontaneous action, perception and reaction.” (http://www.soda.co.uk/explore/eor.htm)

Certainly there is room for more research but it will involve the creation of neural nets that are not of the standard order. If I use them in creating an artificial artist I will have to develop my own approach that redefines their natural usage.

Chapter 5: Proposal

Hofstadter made the humorous suggestion that perhaps AARON could be made to drawings not of people as people do but drawings of drawing machines such as itself. It perhaps could then do a drawing of itself, drawing itself, drawing itself... (Hofstadter Douglas, 1995).

Creating an artificial entity will come down to grasping the full impact of the extent of recursion in human thought. I have yet to gain access to a Lisp application-programming interface but it becomes clear that the logical structure of the code I must generate must take this form. The current platform I am using offers a favourable solution to producing graphics but not the heuristic logic processes that will go into thinking out what the drawing will be.

From my practical research it seemed that neural nets were poorly suited in their application to creating a system parallel to AARON. I do not believe that it is the case as Cohen suggests that neural nets are a means of determining "what is best" (Appendix A), but that the basic models available for study to the initiate of this field are counterproductive to the creative process. They, like noradrenalin, separate articles of information and the faculties of mind as opposed to combining them. But the work of Ed Burton and Stephen Thaler shows that there is an opportunity for the application of neural nets. I could perhaps consider rebuilding the gesture recognition system on a model similar to the creativity machine to create new forms of marks.

Although Cohen shows some resistance to the idea that AARON is rigidly modelled upon a paradigm of human perception we have seen that there are aspects of perception for AARON and us which are very similar. I would argue that nature has something to teach us but I also agree with Kohonen who says that there are inventions of the human mind which excel that of nature, such as the wheel (Kohonen Teuvo, 2001).

From the evidence of Wolfram's automata, it would appear one can use other things than the computer's random number generator. One could achieve that irrationality through the application of some simple iterative rules. Managing that irrationality could be a new source of inspiration for the artificial artist over the random number generator for it would appear that there is an element of repetition as well as randomness.

This could be in fact opened up to the idea of any form of randomness. In London there is a monthly meeting of digital artists who do short talks called Dorkbot: London. At one meeting someone responded to the old chestnut of randomness that perhaps one could use the frequency of London buses as a measurement of randomness in favour of the computer's random number generator.

Perhaps the issue at hand is not necessarily the fact that a system parallel to AARON must have no input - AARON clearly takes input from the random number generator - but that the input must indeed be random. Randomness of input generates the need for intercommunication between the autonomous elements within the system and necessitates the feedback systems to keep that randomness in check. It is the varied quality of information that artists open themselves to that keep their work fresh. By trying new things and seeing new things an artist remains creative.

Bibliography

Arnheim Rudolf, (1954) - 1974 edition, *Art and Visual Perception,* University of California Press

Bishop Christopher M., (1995), *Neural Networks for Pattern Recognition,* Clarendon Press

Bulloch Angela, (1998), *Satellite,* Museum für Gegenwartskunst

Cohen Harold (Cohen Becky, Nil Penny, foreword: Feigenbaum Edward A.), (1984), *The First Artificial Intelligence Coloring Book,* Wiliam Kaufmann, Inc.

Goodman Nelson, (1984), *Language of Art,* Hacket Publishing Company

Greene Rachel, (2004), *Internet Art,* Thames and Hudson

Halsall Francis, (2003), *Art, Art History and Systems – Theory,* University of Glasgow

Hofstader Douglas R., (1979) - 2000 edition, *Gödel, Escher, Bach: and Eternal Golden Braid,* Penguin Books

Hofstadter Douglas R., (1985), *Metamagical Themas,* Penguin books

Hofstadter Douglas R., (1995), *Fluid Concepts and Creative Analogies,* Harper Collins

Hunt V. Daniel, (1985), *Smart Robots,* Chapman and Hall

Kohonen Teuvo, (1977), *Associative Memory,* Springer-Verlag

Kohonen Teuvo, (1989), *Self-Organization and Associative Memory,* Springer - Verlag

Kohonen Teuvo, (2001), *Self-Organizing Maps,* Springer

Livingstone Margaret, (2002), *Vision and Art the Biology of Seeing,* Harry N. Abrams, Inc., Publishers

Maeda John, (2000), *Maeda @ Media,* Thames and Hudson

McCorduck Pamela, (1991), *Aaron's Code: meta-art, artificial intelligence, and the work of Harold Cohen,* Freeman

Morris Frances, (1999), *Chris Burden : When Robots Rule: The Two-Minute Airplane Factory,* Tate Gallery Publishing Ltd.

Raftopoulos Athanassios, (2004), *Cognitive Penetrability of Perception : Attention, Action, Strategies, and Bottom-Up Constraints,* Nova Science Publishers

Tate Gallery, (1983), *Harold Cohen,* Tate Gallery Publications Department

Tyson Keith, (2002), *Head to Hand,* Turnaround

Whitechapel Art Gallery, (May-June 1965), *Harold Cohen : Catalogue of an Exhibition of Paintings 1960-1965,* Whitechapel Art Gallery

Wolfram Stephen, (2002), *A New Kind Of Science,* Wolfram Media Inc.

Internet References

A. Michael Noll's Computer Composition with Lines: http://www.citi.columbia.edu/amnoll/CompArtExamples.html

AARON application download: http://www.kurzweilcyberart.com/

Artificial neuron illustration: http://www.iiit.net/~vikram/nn_intro.html

Back-Propagation neural net: http://www.patol.com/java/NN/index.html

Biological neuron illustration: http://vv.carleton.ca/~neil/neural/neuron-a.html

Ed Burton's EOR: http://www.soda.co.uk/explore/eor.htm

Game of Life example: http://processing.org/learning/examples/cellularautomata1.html

Gesture recognition tutorial: http://www.generation5.org/content/2004/aiTabletOcr.asp

Harold Cohen's papers: http://crca.ucsd.edu/~hcohen/

Hopfield neural net: http://richardbowles.tripod.com/neural/hopfield/hopfield.htm

Kohonen clustering algorithm tutorial: http://www.ai-junkie.com/ann/som/som1.html

Kohonen Self-Organising Map: http://www.patol.com/java/fill/index.html

Mach Bands illustration: http://www.nku.edu/~issues/illusions/MachBands.htm

Stephen Thaler’s Creative Machine: http://www.imagination-engines.com/cm.htm

The Travelling Salesman Problem: http://www.patol.com/java/TSP/index.html

Appendices

Appendix A: Email from Harold Cohen

Dear Aaron Steed

At 11:48 AM 10/30/2005 +0000, you wrote:
Dear Mr. Cohen,
I sent your colleague an email regarding the thesis I am working on and he told me that he had forwarded my request to you.

nothing forwarded to me. who was the colleague?

I thought It would be best for me to continue with my research and see if I could get some basic model together for researching what could be applied to AARON before trying to contact you again.

I've been working on AARON for more than thirty years. Why would you assume I need/want suggestions as to what to apply to AARON, or that I believe AARON needs anything to be applied to it?

There are only four books (exhibition catalogues included) about your work in the British Library catalogue which is why I'm inspired to contact you. The information in them is complete but is perhaps a little dated.
I started learning programming just over a year ago with the ambition to create something like AARON. I was unaware of your work and it was not until I had worked the summer away on the subject that I found I was nearly two-decades late in being innovative, with a machine which has the same name as me. Naturally feeling annoyed I wanted to pick holes and the biggest hole I could find was AARON's seemingly non-adaptive nature

What should it adapt to?

. I spent the following year learning to program and have decided to treat this as a subject for argument.
AARON was described to me as a "top-down" approach to programming. He is essentially the image of Cohen (you) upon the computer. Another top down approach to programming are neural nets, I wondered if neural nets could be introduced to AARON's framework. They would introduce another aspect of yourself upon the system - a biological and adaptive one.

I am interested in knowing if you have applied neural nets at all in AARON's program. If not would you consider it? I would also very much like to know if AARON is chiefly a complex heuristic procedure. Does the program work down through a complex tree of commands to create the image? Lastly I would like to know if work is continuing upon AARON. The last to-do mentioned in the books is the treatment of colour, which has already been dealt with by the downloadable Java demo.

If you type "harold cohen aaron program" to google you'll get about 9000 hits. I'd recommend the web-site of UCSD-research centers-CRCA, where you can find links to about a dozen of my papers dealing with the program.

Yes, work continues, specifically on color. The latest version dispenses with the previous (me-provided) rule-based structure. AARON sets its own goals with respect to color, continuously monitors the developing image and makes its own determination as to how to proceed

in trying to reach its goals.

I didn't notice any concern with color in your java demos.

I've never really understood the top-down vs bottom-up distinction and, no, I've never considered using neural nets, which seem well-suited for matching results already determined to be "best." That doesn't satisfy the condition of art, which is fundamentally an investigative activity, not a production activity.

I don't see why using a computer-centric methodology like neural nets would introduce a biological aspect of myself and I still don't know what you mean by adaptive. By the way, the goal of the project has never been to install "myself" into the program and it has been increasingly directed towards the program's autonomy.

I have found two neural nets so far which looked as if they might offer some value to AARON's program in terms of adaptation.

The Hopfield net could be used to identify favourable patterns chosen by AARON's clients or perhaps similar patten recognition software could suffice (the Hopfield net was pretty unreliable when I expanded its scale from 8x8 to 16x6).

http://www.robotacid.com/PBeta/hopfield2.html

I built a sort of adaptive compostion machine using a Kohonen Self-Organising-Map. My thought was that given an ideal composition AARON could go on to create variants upon this theme. I have made this example as sparse as possible, consider it an abstract treatment of the task - rather like the early versions of AARON. My idea was that the spacer within AARON could become like a sheet of elastic, pulling into areas where there were compositional elements. Further works which please the client would drift into these gravity wells.

Ideal composition? Variants on a theme? Affine/non-affine transformations? Pleasing the client? I begin to suspect that you don't have a clue as to what making art is about. Art is not about -- or rarely about -- playing formal games; eg, the significant difference between "War and Peace" and "Winnie the Pooh" is not to be found by analysis of word frequencies.

http://www.robotacid.com/PBeta/compose

I have enclosed the applets and their source code in a zip file with this email. They can be executed locally using the Processing environment dowloadable from http://processing.org

I realise that these are rather loose and hasty applications to apply to a work as large as AARON. I unfortunately have to guess at the logic you have hammered out over the years. I also realise I have changed course a little from critical analysis to adaptive systems. Having done much more reading I understand it is one thing to encourage a machine to experiment and idealise, quite another to get it to form an opinion.

If you have the time to respond to my mail I would be very grateful. It would great to work on something contributing to an existing cybernetic work of art rather than starting from scratch.

Evidently, you are making some entirely unwarranted assumptions about the availability of AARON for your use. Let me make clear that the code is not available and any attempt to publish "variations" of the program's work as your own would almost certainly be an infringement of copyright.

If you want seriously to work in this area -- as opposed to simply finding a dissertation topic -- I'd recommend that you try to do things yourself, as opposed to applying other people's code to other people's work. For that you'll need a lot more than two years of programming in a suitable language -- lisp would be the obvious choice -- and a great deal more understanding of what art-making IS than you appear to have.

I find myself wondering why, of all the requests for information I get from graduate students all over the world, I would find yours (alone!) to be offensive. I suppose Cezanne might have found it similarly offensive if a beginner assumed he could "improve" on his paintings for him.

Harold Cohen

Regards,

Aaron Steed.

Appendix B: Email from Douglas Hofstadter

Dear Mr. Steed,

Thanks for your note and the kind words about my writings.
I didn't totally understand the Wolfram quote, but in any case here's a reaction.
There's no doubt that simple initial conditions, if they are fed into certain kinds of rather simple iterative rules (and keep in mind that an iterative rule is all that any cellular automaton ever is), can produce very random-looking, wild, unpredictable output. In fact, Wolfram uses my own recursive function Q(n) somewhere in his book as an example of exactly that phenomenon. From a trivial input -- namely, two 1's, that's all -- and a rather simple nested recursion -- Q(n) = Q(n-Q(n-1)) + Q(n-Q(n-2)) -- then presto! you get wild, crazy, random-seeming behavior. So Wolfram's idea is hardly news to me -- I discovered this fact for myself in the early 1960's. (That's when I invented the Q function, studied it in detail, and discovered its fascinating chaotic properties.) But I don't see what this fact has to do with fluids or fluidity.
Physical fluidity is a statistically robust kind of behavior that results from huge numbers of infinitesimal specks constantly bashing against each other randomly, and only from vast numbers of such random bashings can you get the emergent fluid behavior that we are familiar with (for example, the behavior we see when we pour a pitcher of milk onto a horizontal table). If, in a computer simulation of a fluid, you replace the (pseudo-)molecules' randomness by pseudo-randomness, that'll still work, as long as the pseudo-random numbers are good enough (i.e., if they pass certain key statistical tests, somewhat the way the digits of pi seem very random for most purposes, even though they are the furthest thing in the world from being truly random). And even if the pseudo-random numbers are generated by a very simple iteration or recursion (as they almost always are, in computers), such as my Q function or some cellular automaton or even the digits of pi, that's fine -- you'll still get fluidity as the emergent outcome.
In short, Wolfram's quote, at least to the extent that I understand it, doesn't in the least undermine my claim about fluidity being emergent. It seems to me that he's merely claiming that a very simple cellular automaton (or some other type of simple number-theoretical function) can produce high-quality pseudo-random numbers that can then be used in simulating the bashings-together of simulated molecules to produce very realistic-seeming emergent fluidity. Fine! Terrific! I never would have said otherwise. That's the whole idea behind the use of pseudo-random numbers in Copycat.
In short, the statements made in "FCCA" about randomness and emergent fluidity (and thus, indirectly, about creativity) are perhaps "old" (I don't think of ten years as being old, but that's another matter), but they seem to me to be as valid as ever.
I hope this helps you out a bit. Best wishes!

Sincerely,
Douglas Hofstadter.