Computational modelling of creativity in abstract art.

Mike Holcombe, Samantha Smith, Rowan Merewood and Andy Swingeford.

Department of Computer Science,
Univesity of Sheffield,
Portobello Street,
SHEFFIELD, S1 4DP,
{m.holcombe@dcs.shef.ac.uk}

***Abstract**: Artistic creativity is studied through the construction of computational models of a number of well known modern artists. In particular, the work of Piet Modrian, M.C. Escher and Paul Klee are suitable vehicles for investigation since their work is accompanied by extensive writings describing the ideas and motivation behind their compositions. In particular, we have tried to abstract from their theories rules that describe the construction procsss or the properties that their finished artefacts possess in order to create software programmes that can articulate these rules.*
In this way we are able to simulate either automatically or with user interaction, the process of creating works of art of a similar genre and satisfying the properties desired by the artist. Since the rules are bound to be considerably more complex than those currently exposed we are looking to use machine learning techniques to develop more sophisticated agents whihc may behave more closely like the actual artist.

Keywords: Software agent, cognitive model, computational model, artistic creativity, visual art.

1. Introduction.

A number of researchers have attempted to provide software environments that allow users to create artistic images in the style of well known artists. Perhaps the most sophisticated such system is *Aaron* [COHE95] which is a large programme that captures the stylistic techniques of a number of artists and styles. Such a system requires the user to create a concept and composition but provides a suite of supporting functions that can be exploited in developing the picture according to the pallettes and techniques of specific artists. The system demonstrates the power of developing virtual art and the application of “virtual reality” to the visual arts.

It may be possible to examine the creative process at a more fundamental level using this type of approach but we have embarked on a different strategy in order to understand these issues. We are interested in trying to encapsulate creativity, as far as is possible, within the context of a computational problem by building *computation models* of the artistic process.

Our basic assumption is that the process of painting a picture can be mapped onto a computational process. This raises many issues of a philosphical nature which we will not discuss further here, however, we have had some success in this direction and that is what is reported in this paper. It is certainly possible, in some cases, to create a software system that will generate art works automatically as well as with some user intervention, in the style of a particular painter and this is reported on below.

2. The computational framework.

Software agents are autonomous programs which will behave according to certain prescribed and built in rules, reacting to their environment and, in some cases, learning from their experience in order to adapt their behaviour so as to satisfy their objectives. They can exist solely in an environment or there could be a community of agents interacting with each other as well as with their environment.

Each agent, since it is a computer programme, can be thought of as a computational model. We have devised a formal, rigorous notation for specifying such agents either singly or as a network of communicating agents. This provides us with a mechanism for analysing theoretically the behaviour of such systems. [KEF02]

3. Mondrian.

Through many years of experimentation and evolution, Piet Mondrian developed a unique and very distinctive painting style that has influenced everything from advertising to architecture. This style was based entirely on a set of theories about what Mondrian wanted his art to portray and how he wanted to represent the world in his paintings. The main theory being that the images he painted should create a pure reality, which in its simplest form means reducing structures and forms to their most basic components, i.e. straight black lines, and reducing colour to rectangular blocks of primary colour (red, yellow and blue) or non-colours (white, black and grey), from which it would be possible to create any other colour. This expression of style in its most basic form, meant that Mondrian had to develop new techniques to express what his work represented. Working with the fields of symmetry, equilibrium, and a varying complexity of shape and form derived from the static structure of the lines. However, despite the fact that Mondrian used the most basic of structures, he was still able to express motion, rotation, strength, speed and many other aspects that might not be considered when an individual encounters the work of Mondrian for the first time. The majority of Mondrian s theoretical writings were published in De Stijl between 1917 and 1924, however most fail to relate directly to his art and are at best difficult to comprehend. This poses a great problem, for without the theory or thinking behind Mondrian's work, how will it be possible to model those processes in an intelligent piece of software? The solution seemed not to concentrate solely on the writings of Mondrian but instead to analyse his paintings, understand the concepts behind the painting and encode these concepts or rules for a feasible Mondrian image into the system.

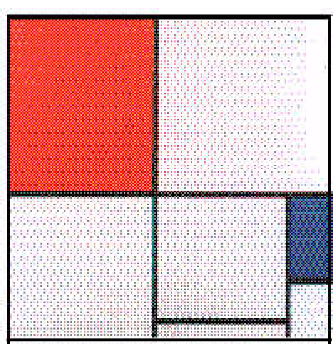

Composition A

3.1 The Decomposition of a Mondrian Painting Every aspect of a Mondrian painting was carefully planned so as to produce the desired impression. Therefore to artificially create an image in the style of Mondrian it is essential to be able to understand the artists original intentions. The following extract describes in detail the purpose of each aspect of Composition with Red, Yellow and Blue, 1921 the diagrammatic form of which is shown below.

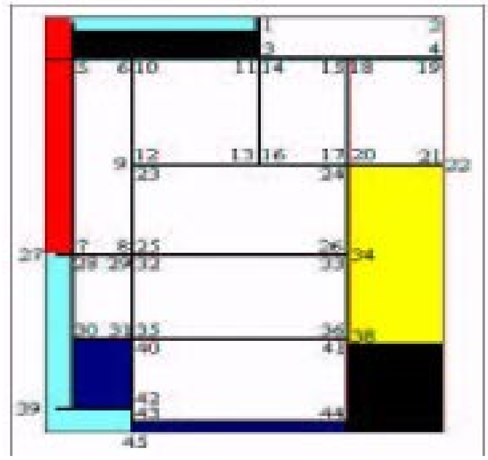

Mondrian Composition with Red, Yellow and Blue, 1921

The long corridor which starts at the base from points 45 & 46 and extends up to 10 & 15 at the top invites perception of the lines 45-10 and 46-15 as boundaries. However, these boundary lines are intersected by horizontals twice on the left and once on the right. The most active lines here are 27-34, which cuts right across the rectangle 23-35-36-24, and 9-22 which cuts through the rectangle 10-25-26-15, both of which, if they had remained intact could have functioned as stable focal points for the composition... Moreover, the horizontal line 9-22 itself forms the base of three further rectangles. A left-right directionality is created by the fact that the vertical 11-13 is positioned right of centre of the rectangle 10-12-17-15, creating an asymmetrical effect which is in turn balanced by the rectangle 14-16-21-19. This rectangle is further subdivided, creating an irresolvable tension between movement towards the inside and towards the outside of the picture. The rectangle 10-12-17-15 cannot function as a stable focal point, because it is destabilised by 14-16-17-15 (which is white), part of 10-25-26-15. It therefore directs the viewer s gaze into and down the central corridor , whereas 18-20-21-19, which is both open and bluish white, like 1-3-4-2 above it, directs the viewers gaze out of the picture. The predominating direction of this outward movement is upwards and to the right, accentuated by the fact that the three lines in the lower left stop just short of the edge, directing the spectator s gaze inwards and upwards. In this way the upper right and self-reflexive grid becomes expansive. [REYN95] It is clear from the breakdown of Composition with Red, Yellow and Blue , that some of Mondrian s paintings were incredibly complicated and detailed works of art which may be impossible to model due to their complexity. This dissertation will therefore focus on the simplified problem of analysing and trying to produce images of the complexity and style of Composition A. This allows for a significant reduction in the complexity of the problem addressed but still allows for expansion by the addition of subsequent rules once basic rules have been established and the software is working as expected.

3.2 Classification of Mondrian Images

For the purpose of this paper, a set of five of Mondrian's paintings have been divided into five different genres. This is designed to enable the creation of a system that works towards certain goals depending on the genre that it is trying to recreate instead of creating a random image of lines and colours. It also provides a facility for user-input, allowing the user to specify the style of image to be created. The examples and explanations of the 5 genres are defined as follows:

Conflict: The painting appears to be divided into two distinct halves with a vertical line running down the centre with the left side of the picture seeming to pull upwards and the right side pulling down. Key Points ~ Divide the image, reflect and mirror the effect that one side of the image has so that the image is instantly put into conflict.

Calm: The large expanse of white in the centre of the painting provides a calming focal point for the viewer and even though the colours around the edge of the painting try to rotate, they are stopped by the block of red in the bottom right. So removing motion from the picture. Key Points ~ The image must have an uncluttered structure and contain no motion.

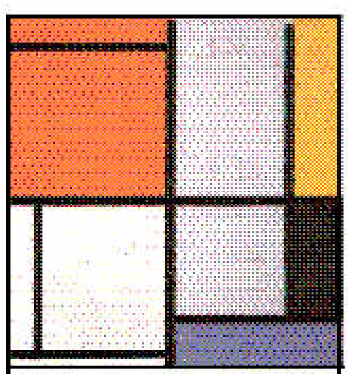

Conflict

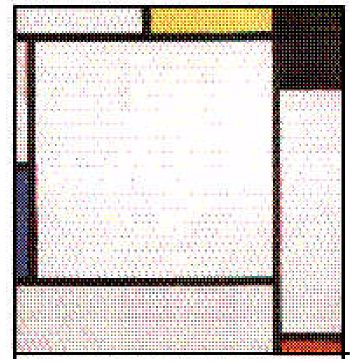

Calm

Architecture: The use of strong, structured lines, combined with a style of colouring that suggests the creation of layering or building from bottom to top gives the entire picture a look more usually associated with a plan than with a painting, so defining the clear architectural style. Key Points ~ Clear use of an axis to provide height and width to the picture. The use of strong, regular blocks produces the stacking effect, the image must not be overworked so as to destroy the simple structure.

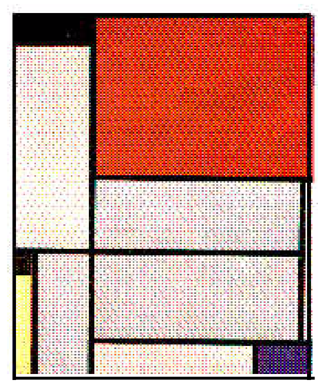

Architecture

Rotation: The painting uses an off-set axis to provide the focal point for the eye and at the same time the point of rotation for the image. Helping the feeling of rotation, the blocks of colour are well balanced and seem to be pushing on the open expanses of white, so as to increase the feeling of rotation. Key Points ~ There must be an axis or turning point and the different aspects of the painting must be balanced using reflection or inversion so that equilibrium is formed that could be imagined to be rotating.

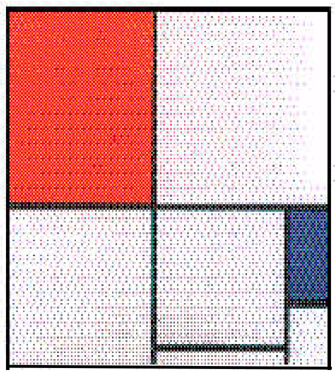

Rotation

Parallelism: The painting uses a technical quality of mechanical repetition which reinforces the abstract character of the picture [DEICH95]. The only abstraction from this is the artistic impression that seems to creep into the picture in the form of the coloured blocks. Key Points ~ The image must be constructed using lines that maintain a parallel connection to the rest of the picture and extend to intersect with another line or the edge of the image.

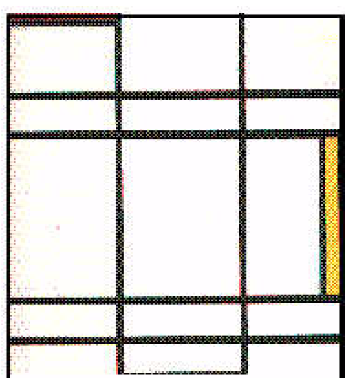

Parallelism

3.3 The Rules of Neo-Plastic Art

Using an essay written in 1927 (the dwelling-the street-the town) [MOND93], it is possible to state six specific rules or requirements of Neo-Plastic art, all of which are adhered to by Mondrian's pictures. These will form the basis for all computational logic or reasoning used by the Mondrian generator to construct individual Mondrian images. The Neo-Plastic rules, stated by Mondrian are as follows: RULE 1 The plastic means must be the rectangular plane or prism in primary colours (red, blue and yellow) and in non-colour (white, black and grey). In architecture, empty space can be counted as non-colour, denaturalised material as colour.

RULE 2 Equivalence in the dimension and colour of the plastic means is necessary. Although varying in dimension and colour, the plastic means will nevertheless have an equal value. Generally, equilibrium implies a large area of non-colour or empty space opposed to a comparatively small area of colour or material.

RULE 3 Just as dual opposition is required in the plastic means, it is also required in the composition.

RULE 4 Constant equilibrium is achieved by the relationship of position and is expressed by the straight line (boundary of the pure plastic means) in its principle, perpendicular, opposition.

RULE 5 Equilibrium that neutralises and annihilates the plastic means is achieved through the relationship of proportion in which they are placed and which create a vital rhythm.

RULE 6 Naturalistic repetition, symmetry, must be excluded.

Here, then, are six Neo-Plastic laws that determine the pure plastic means and how they are used. The use of these rules, combined with rules gained from the analysis of Mondrian s paintings (section 3.2) should enable the creation of software that adheres to the principles of Neo-Plasticism and produces art in the same style as Mondrian.

3.4. *Generic Rules for all Mondrian Images* Using the rules shown above it is possible to formalise the criteria of a Mondrian image as follows: *Colour*: The image must only feature the colours red, yellow, blue, grey, black & white. *Shape*: The image may only contain rectangular shaped planes and lines. If two lines run in the same plane, they are parallel. If two lines intersect, the angle of intersection = 90˚. By specifying these two criteria, an image can only be made from intersecting lines and may only contain rectangular shapes. *Colour:Non-Color Ratio*. Colour has a dominant effect over non-colour. Therefore There should never be more colour in a picture than non-colour. In general, a minimum ration of 2:1 non-colour to colour. This may vary depending on the rules of the specific style of image created.

3.5 *Relationship of Complexity to Simplicity*: There must be a co-ordination between areas of complexity and areas of empty space. The two must complement each other rather than compete for space. Therefore areas of complexity and simplicity must be placed in opposite corners of an image to create balance. Symmetry: No Mondrian image may contain symmetry.

3.6 *Specific Rules for Mondrian Images* The generic rules stated above must hold for any Mondrian image, however there are also specific rules that must be applied, depending on the style of image created. These rules are formed from the analysis of Mondrian's different styles as carried out in section 3.2 and are intended to express basic algorithms for the production of each type of Mondrian image. 3.6.1 Conflict Image Creation: 1) Subdivide the canvas, horizontally roughly in the centre. 2) Subdivide the canvas vertically to form a central cross. 3) Choose a quadrant. 4) Subdivide the quadrant and its diagonally opposite quadrant. 5) Sparsely colour the two quadrants. 6) Subdivide the remaining quadrants. 7) Intensely colour the remaining quadrants. This produces the opposition within the image and generates the effect of one side of the image pulling away from the other, i.e. conflict between sides.

3.6.2 Calm Image Creation: 1) Subdivide the canvas horizontally at the top and bottom. 2) Subdivide the canvas vertically, left and right so as to form a large focal square of empty space. 3) Subdivide complete lines, maintaining the simplistic nature of the image. Never invade the focal square with any line. 4) Fill 3-4 areas with colour, never filling the focal square. 3.6.3 Architecture Image Creation: 1) Create an axis, vertically offset from centre. 2) Add a large area of colour to the larger side of the offset. 3) Below the area of colour, horizontally subdivide the remaining space into approximately equal blocks. 4) Minimally subdivide the sections created with the use of coloured blocks.

3.6.4 Rotation Image Creation: 1) Subdivide the canvas, horizontally roughly in the centre. 2) Subdivide the canvas vertically to form a central cross. 3) Choose a quadrant of the image. 4) Colour chosen quadrant. 5) Subdivide diagonally opposite quadrant vertically. 6) Subdivide diagonally opposite quadrant horizontally twice. 7) Fill one of the sub-divisions with a colour different to the colour of the opposite quadrant.

3.6.5 Parallelism Image Creation: 1) Divide the canvas vertically using two parallel lines. 2) Divide the canvas horizontally using parallel lines with equal spacing between pairs of lines. 3) Subdivide the compartments created with a minimal use of additional lines. 4) Fill a maximum of three small compartments with colour.

As can be seen from the screen shots, the interface includes a pull-down menu box, style control sliders for mood setting and a window provides a complete set of image manipulation tools for the user, allowing the user to edit every aspect of an image once the image has been created by the software.

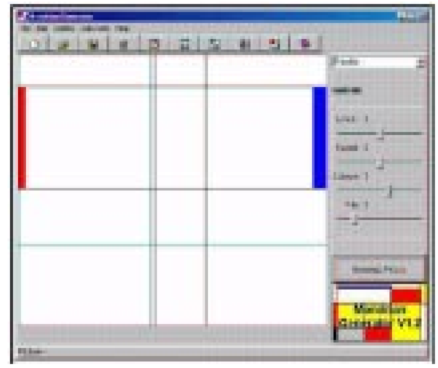

Tool Screenshot

This tool succeeded in implementing a number of the rules of Neo-Plastic art. Five individual styles of image were implemented, each of which used image specific criteria when producing an image. These criteria were not obtainable by simple extraction from the writings of Mondrian. Instead, they were interpreted from analysis of his paintings and by studying the decompilation of his painting style as discussed by critics in books concerning Mondrian [REYN95].

4. Escher.

The Dutch graphic artist M.C. Escher was fascinated by the 'art' of regular plane filling, in particular with the goal of filling a finite plane in such a way that it appears to stretch on for infinity. As a pioneer in his field Escher looked to other cultures to find similar (but more primitive) plane filling techniques including Moorish and Japanese artists. However much of Escher's work was a result of his own experimentation with plane filling. Fortunately Escher chose to write down a great deal of what he had learned over a lifetime, probably spurred on by the knowledge that certainly up until that time he was the only individual working in that field.

4.1 Background.

Escher attempted to portray infinity in several ways, for example through tessellation, by depicting scenes that show a process repeating over and over again and by using curved surfaces. Escher's more complex and intricate work where characters are shown ascending or descending an impossible staircase are almost certainly beyond the scope of this project, and so it will be necessary to restrict it to Escher's simpler, but by no means simple tessellation work.

Symmetry work 96, 1955

The intention of this research is to develop a piece of software that can emulate Escher's style of tessellated plane filling. In order to do this it will be necessary to extract and use the rules and guidelines that Escher worked to, in order to develop software that can emulate his methods and hopefully also his style.

4.2. Objectives.

In order to achieve our main objective, several sub goals must be achieved first.

First it is necessary to extract Escher's rules and guidelines from his own written documentation of his work in the book Escher on Escher, [ESCH1989]. It is also important to consider and attempt to record in written form what separates an Escher design from a random tessellation. In effect to capture what *defines* Escher's work. The conclusions drawn from extracting Escher's rules and in capturing what defines Escher's work can hopefully then be used to define the criteria that would specify a successful art agent for the work of M.C. Escher.

His procedure was to create a simple tesselation of the plane using parallelograms and then to perturb their boundaries in a way that created two complementary shapes. He then added detail as shown.

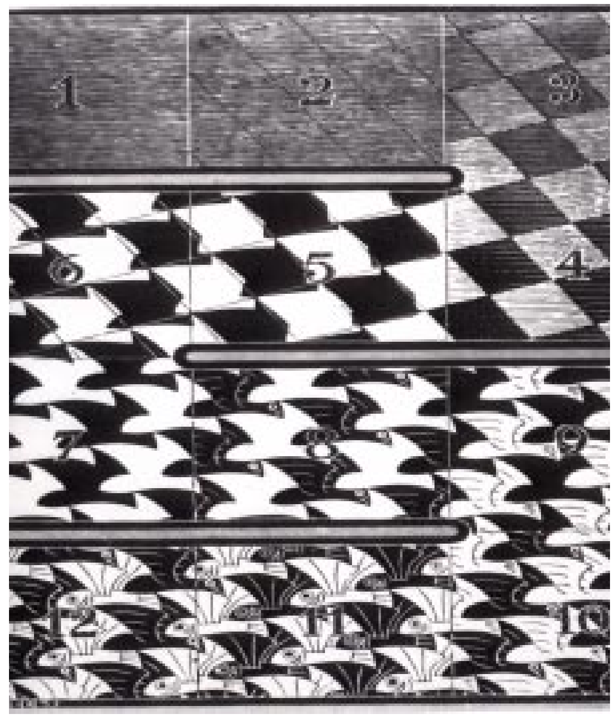

Escher's procedure illustrated

4.3 The Escher tool.

Once Escher's work has been analysed as above it should then be possible to begin to translate his work into a system specification that comprises algorithms and constraints that can be incorporated into a piece of software. As well as establishing the requirements that the software must meet in order to emulate Escher's plane filling work it is also necessary to consider the needs and requirements of a human end user in interacting with the software.

Over his career Escher produced a large and diverse collection of Tessellations. They vary in style but we concentrate, here on 2 dimensional examples. The software tool built to emulate this is shown below where the basic tesselation is established and is being manipulated to create complementary and contrasting images prior to embellishment.

.

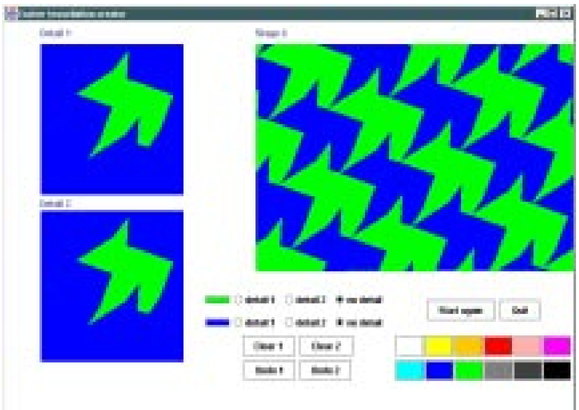

An example of an image constructed with the tool with its genesis is:

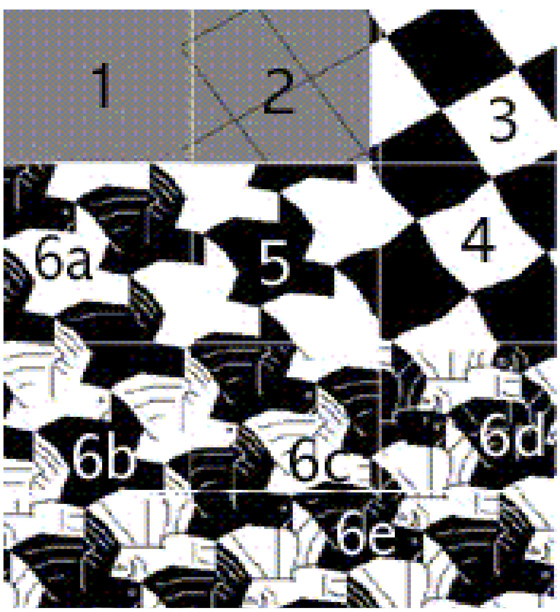

Another image from the tool:

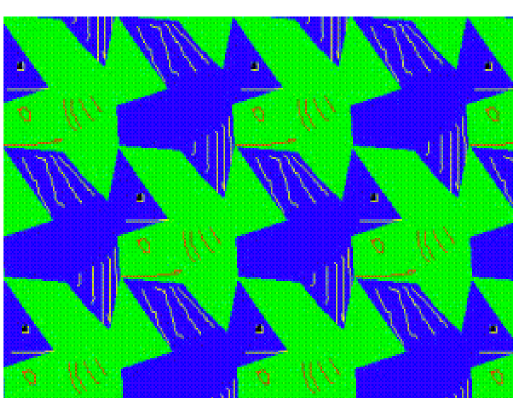

5. Klee.

In a similar vein, Paul Klee, also expressed his method of working in an almost mathematical fashion.[KLEE68]. If these concepts can be distilled into programatical rules then they may be applied according to a set of criteria to generate an image.

5.1 Background.

The Pedagogical Sketchbook [KLEE68] was written by Klee during his time at the Bauhaus. It provided the basis for one of the courses that Klee taught and represents an e ort to analytically describe the process of artistic creation with almost mathematical precision. It begins with the basic building block of an image, the line, and progresses to moving onto proportion, structure, balance, gravity and ending on kinetic and chromatic energy within the image. For each section, Klee relates methods to particular feelings they evoke or an aspect of nature being represented. The continual linking of the image to natural surroundings should be noted, as it poses a difficult question for the program. How can the knowledge of our natural environment be imparted to the agents within the system? However, if we focus merely on the techniques then the appeal of Klee's descriptions are in the linear, almost modular approach to application. An excellent example can be seen in his illustration of an active line with complementary forms. That is (to paraphrase) a single line accompanied by multiple tangential lines, a line circumscribing itself or two secondary lines moving around a single, imaginary line.

It is this modular approach which lays out an obvious route to writing agents to perform these tasks, which will be expanded on later.

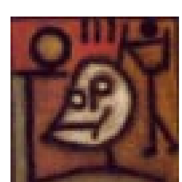
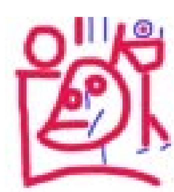

Tod und feuer (Death and Fire) 1940

The works focused on in this paper come mainly from the latter part of Klee's life. As his disease became worse, so his style evolved to the broad, powerful brush stokes and the large blocks of colour. From a purely technical point of view, this provides an easier analysis than the more detailed or "biological" works, definite edges and separations allow clearer identi cation of pictorial elements. It should also be noted that the structure illustrated in these studies is just one interpretation of the image; art is not a science and it should not be assumed that any experiment is repeatable to the degree of accuracy demanded in an engineering discipline. What has not been addressed in the analysis of these images is their meaning, or how they actually qualify as art. Unfortunately, the implementation of a knowledge base (or similar construct) to allow the application to reason about the artistic merit of an image is beyond the scope of this project. However, it would be an interesting to see the reactions to the out- put of a program that operated in the way described compared to the solely rule-based foundation of the application developed here.

5.2. Analysis of some of Klee's images.

Umgriff, (Hold), 1939

Since this painting features no real colouring, it makes a good starting point to examine only the structural aspect of the piece. The background may be regarded as white or empty space into which the artist places the pictorial elements. In this case, the elements are represented by the red in the diagram. There are only three types of basic element here: a straight line, a curved line and a point. None of these cross each other, but they can join. There are also points where they run roughly parallel to each other. The blue represents non-structural additions to the image, embellishing the existing elements. The top horizontal line is complemented by two circles alongside its lower edge. The left-most curve also has a circle in close proximity. Our point is enlarged to become a solid circle, but is structurally still only a point. Finally, the vertical line in the right corner has a horizontal bar added to one end.

The analysis of Klee's images have been broken down into two stage developments; the generation of the base elements followed by a variety of embellishments on top of them. However, what has not been covered is how the elements were arranged spatially in the first place. At this point, the compositional rules may be used as foundation for the placement of these first pieces. The creation of the image, by these rules, is then a three tier process. This is the process which the application must implement.

5.3 Agent architecture and process model.

The agents are the only components in the system that can affect the contents of the image. For an agent to do this, it must first determine whether its operation will fit with the current theme of the image as it stands. The actual mechanism for this will be covered later, the important concept here is that there is a communication between the agent and the higher-level section of the system before any alteration of the image takes place. If an agent is to act on the image, then its initiation function is called, with a reference to the subsection (although, this may well be the entire image) of the piece for it to modify. It may also be possible that the agent returns some sort of message to indicate the success of failure of its actions. Since an agent-based approach is somewhat analogous to the object- orientated paradigm, an agent can be directly represented by a Java class.

This eliminates the need to include a separate language for the agents within the system, and also allows access to Java's extensive API for their operation. Implementing an agent is then a case of creating a class that inherits from the agent superclass of the relevant tier.

Since agents are a dynamic aspect of the system their instantiation may not be hard-coded into the application as is possible with static elements. If this was the case, then the system would have to be re-compiled with the addition of each new agent, certainly impractical if anyone other than the developer wishes to add their own agents. The agent pool creates an instantiation of each agent for its tier. This is achieved by keeping all agents in a specific directory and following a naming convention for each tier. Each agent has a two-letter prefix to identify itself to the agent pool: "SA" for the skeletal layer, "MA" for the muscle layer and "KA" for the final skin layer. A method in the pool then searches a given directory for all files matching the prefix *. class pattern and creates an instance of each, stored in a Vector.

The agent pool also serves to abstract the details of file handling and instantiation away from the higher manager layer. The collection of agents is private to the pool, access being given through a set of acessor methods. The manager sits at the top of the system and handles interactions with any external processes, such as the user interface. Therefore, the manager must provide the necessary access to the agents and the agent pool that are below it.

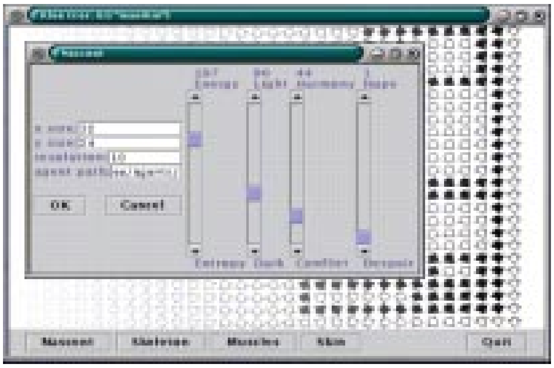

A composite screen shot showing the mood bars.

The user should be able to perform the tasks, listed in approximately the order in which they would occur: Set dimensions of the Skeletal matrix; Set resolution of the final image; Set the mood of the image; Manipulate the set of Skeletal agents to be used; Generate the Skeletal layer; Display a visualisation of the Skeletal layer;

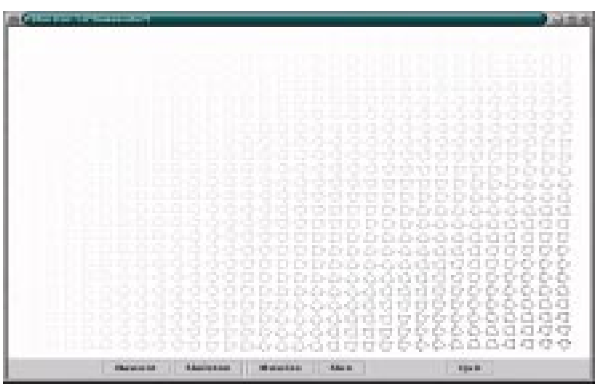

Skeletal layer for *fish*

Manipulate the set of Muscle agents to be used; Generate the Muscle layer; Display a visualisation of the Muscle layer;

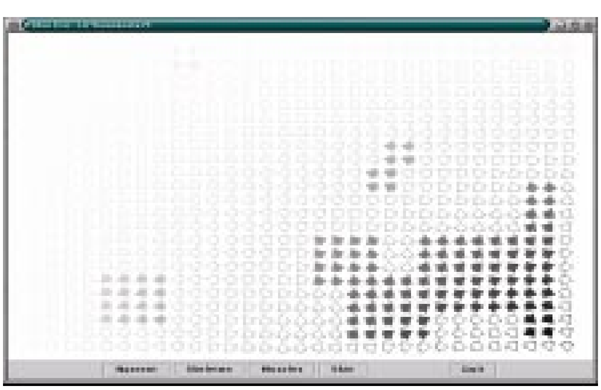

Muscle layer for *fish*

Manipulate the set of Skin agents to be used; Generate the Skin layer

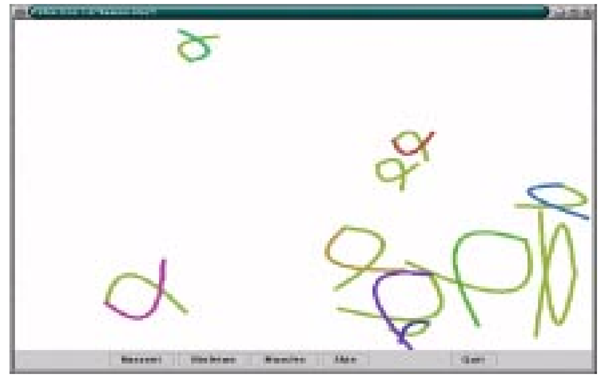

Skin layer for *fish*

Currently the tool is a framework which allows for the addition of large numbers of agents to carry out drawing activities. These agents are simple java classes and can be added at will. These pictures were drawn using a simple abstract fish like image. With a much larger pool of agents much more interesting images will be possible.

5. Conclusions.

The aim of this project was to investigate the possibility of taking some examples of abstract art and generate a software tool that would allow a user to create similar work. After investigating the various styles of work, Mondrian's mood compositions and their rules, Escher's mode of regular plane filling using translation and Klee's modular and layered approach to composition we have built software systems to emulate these artistic processes.

There are many possible extensions to this work. Those extensions may entail further investigation into allowing the creation of more of these artists' modes of work, or applying the same ideas to a different artists in order to extract the necessary information to recreate some of their work in a similar way. It is unlikely that all forms of art will lend themselves to this form of analysis. Artists that seem to follow rules of any kind are probably the best candidates.

In order to refine te approach further we would like to use types of machine learning to train the tools - essentially their underlying rules - by comparing the output from the tools with the *real thing* and evolving the systems further towards real art generators.

In concluding we quote the comments of the Diana Syder, Poet in residence of the Deparmten of Electronic and Electrica Engineering at the University of Sheffield who saw these tools demonstrated.

“the generating art [software] raises all sorts of questions about what art is, and blurs the edges of that...is art it in the mind of the executor or the receiver...how much is the person who sees the piece the artist (ie has the work of putting their own meanings on to a given piece of art) c.f. the person who made it...and how important is the process, as in the journey the artist makes to get to an end point.”

6. References.

[ESCH1989] Escher, M.C. (1989) "Escher on Escher Exploring the Infinite" Abrams

[ESCH1989A] Escher, M.C. (1989) "The Graphic Work" Taschen

[ESCH1985] Escher, M.C. (1995) "Postcardbook M.C Escher" Taschen

[HOFS1979] Hofstander, Douglas F (1979) "Gödel, Escher, Bach: An eternal Golden Braid" Penguin Books

[SWIN2001] Swingewood, Andrew (2001) "Intelligent Agents and Artistic Creation" Department of Computer Science, University of Sheffield.

[SPAL2002] Spalter, Anne Morgan "The computer in the visual arts". Addison Wesley

www.groups.dcs.st.and.ac.uk/~history/Mathematicians/Escher.html Found November 2001

[COHE79] Harold Cohen, “What is an Image?"Proceedings of the Interna- tional Conference on Arti cial Intelligence (IJCAI), 1979

[COHE88] Harold Cohen, “How to Draw Three People in a Botanical Garden", Proceedings of the Seventh National Conference on Arti cial Intelligence, pps.846-855

[COHE95] Harold Cohen, “The Further Exploits of Aaron, Painter", Stanford Humanities Review, Volume 4, Issue 2 : Constructions of the Mind, 1995

[DARP93] The DARPA Knowledge Sharing Initiative, ‡RAFT Speci - cation of the KQML Agent-Communication Language", The DARPA Knowledge Sharing Initiative: External Interfaces Working Group, 1993

[FELD99] Susan Feldman, Edmund Yu, “Intelligent Agents: A Primer", Searcher, Volume 7, Issue 9, 1999

[GILB95] Don Gilbert, Manny Aparicio, Betty Atkinson, Steve Brady, Joe Ciccarino, Benjamin Grosof, Pat O’Connor, Damian Osisek, Steve Pritko, Rick Spagna, and Les Wilson, IBM Intelligent Agent Strategy, white paper, 1995.

[HALL92] Douglas Hall, “Klee", 2nd Edition, Phaidon Press Limited, 1992

[JENN94] Michael Woolridge, Nick Jennings, “Intelligent Agents: Theory and Practice", Submitted to: Knowledge Engineering Review, 1994

[KLEE68] Paul Klee, “Pedagogical Sketchbook", Faber and Faber Ltd., 1968 (originally Paul Klee, “Padagogisches Skizzenbuch", Bauhaus Books, 1925)

[KLEE69] Paul Klee, "On Modern Art", Faber and Faber Ltd., 1969

[KUDI02] Robert Kudielka, Bridget Riley, “Paul Klee: The Nature of Cre- ation", Hayward Gallery Publishing, 2002

[KURZ01] Kurzweil CyberArt Technologies, Inc., “Aaron, The Cybernetic Artist", http://www.kurzweilcyberart.com/, 2001

[MERE02] Rowan Merewood, “On-line javadoc Documentation", http://www.nekrodomos.clara.net/clockwise/dissertation /javadoc/ [MERE02] Rowan Merewood, “Intelligent agents and artistic creation”, Department of Computer Science report, University of Sheffield, 2002

[NEWT87] Isaac Newton, \Philosophiae Naturalis Principia Mathematica", Volume 1, Cambridge University Press, 1687

[SMITH02] Samantah Smith, “Art emulation", Department of Computer Science, University of She eld, 2002

[TURI50] Alan Turing, “Computing, machinery and intelligence", Mind, Volume 59, Issue 236, pps. 433-60, 1950

[PIET26] Piet Mondrian, The dwelling the street the city. This essay first appeared in the first number of i 10 under the title Neo-Plasticism: De woning de straat de stad , 1927. [KEIL89] F.C. Keil, Concepts, Kinds, and Cognitive Development. Communications of the ACM, 37, 7, 143-146, 1989.

[MOND93] Piet Mondrian (edited & translated by Holtzman and James), The New Art The New Life. The Collected Writings of Piet Mondrian. Da Capo Press. New York, 1993. [LOCH94] Hans Locher, Piet Mondrian: Colour, Structure and Symbolism. Gachnang & Springer, 1994.

[WARN94] Carsten - Peter Warncke, De Stijl 1917-1931. Benedikt Taschen, 1994. [DEICH95] Susan Deicher, Mondrian 1872-1944: Structures in Space. Benedikt Taschen, 1995.

[REYN95] Dee Reynolds, Symbolist Aesthetics and Early Abstract Art. Cambridge University Press, 1995.

[FRAN96] Stan Franklin & Art Graesser, Is it an Agent, or Just a Program?: A Taxonomy for Autonomous Agents. Third International Workshop on Agent Theories, Architectures and Languages, Springer-Verlag, 1996. [RUSS95] Stuart Russell & Peter Norvig, Artificial Intelligence: A Modern Approach. Prentice Hall, 1995.

[KEF02] Petros Kefalas, Mike Holcombe & Marian Gheorghe, “Formal specification of agents”, to appear.