TITLE WET AND STICKY: A NOVEL MODEL FOR COMPUTER-BASED PAINTING

AUTHOR Malcolm Tunde COCKSHOTT

DEGREE Ph.D

AWARDING BODY Glasgow University

DATE 1991

THESIS NUMBER DX205420

WET AND STICKY: A NOVEL MODEL FOR COMPUTER-BASED PAINTING.

by

Malcolm Tunde Cockshott, B.A.Hons.

being a thesis submitted for the degree of Doctor of Philosophy in the Department of Computing Science at the University of Glasgow.

13 June 1991.

NUMEROUS ORIGINALS
IN COLOUR

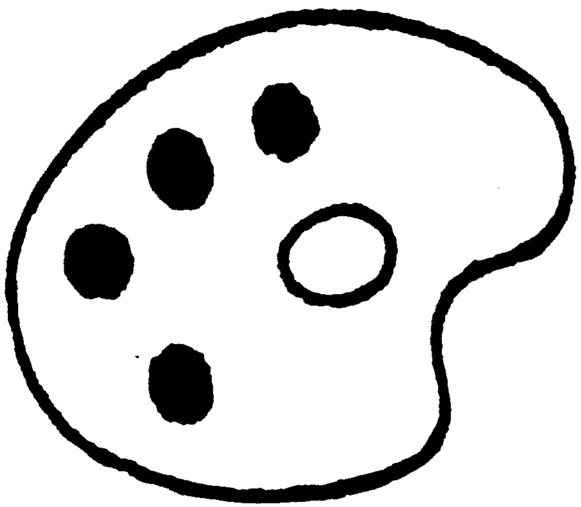

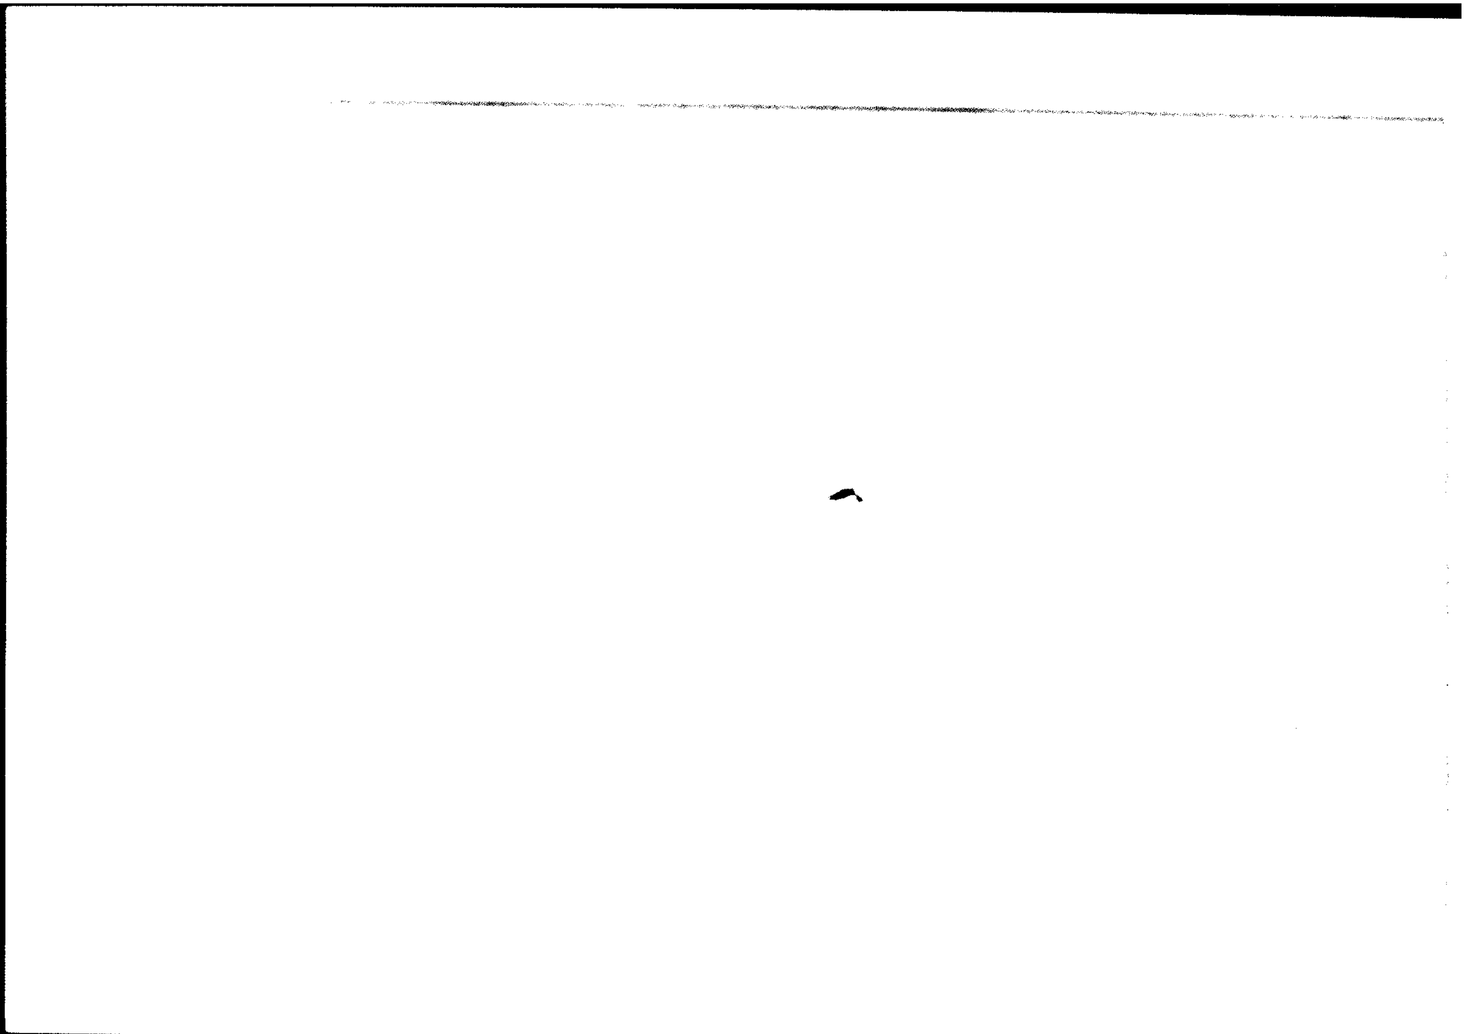

An example of the mark making capabilities of the Wet and Sticky model.
In this case an image of a bull was digitised for computer use, and then Wet and Sticky paint was applied. In addition the left-hand half of the image was given a gravity direction of up, relative to the image, and the right hand side was given a gravi

Table of Contents

Table of Figures

List of Colour Plates

Acknowledgments.

In acknowledging the support I have received over the period of my thesis it is hard to mention everyone involved, but there are a few individuals without who's support the tasks of research and writing–up would have been very hard. In particular I would like to thank my wife Kay for all her support and badgering encouragement, my brother Paul for his advice and interest, and all the members of the department who helped in so many ways. My special thanks to Kevin Waite, and David England for their constant support and advice with programming problems, and to Naveed Khan for his apparent incredulity that my work would ever finish. Finally I would like to thank my supervisor, John Patterson, for his help in turning a computer illiterate sculptor into a PhD candidate.

Financial support for this work was provided in the form of a Scientific and Engineering Research Council research studentship.

Declaration of Originality.

The material presented in this thesis is entirely the result of my own independent research carried out in the Department of Computing Science of the University of Glasgow, under the supervision of Dr John Patterson. Any published or un-published material that is used by me has been given full acknowledgment in the text.

Permission to Copy

Abbreviations etc.

After the first mention of the Wet and Sticky model it appears in the format WET & STICKY. This serves as an aid to legibility and to emphasis. Within this thesis the use of the word *he* as a reference to a user etc. is an abbrieviation for he/she.

Summary

The problems of computer based painting are considered from a fine arts stand point. A detailed criticism of existing approaches is provided. This criticism centres on the limited depth of existing models and the resultant limited complexity and relative poverty of their mark making capabilities. The notion of the importance of the accidental in art is highlighted and an argument is made for its inclusion in computer based painting models. An informal task analysis is carried out and a description is provided for the task of domain modeling. The results of this task analysis confirm the inadequacy of the existing computer based painting models.

A novel paradigm, Wet & Sticky, is proposed which models the physical and behavioural characteristics of paint rather than just its colour properties. The initial proposals for the model require that it mimics the actions of gravity and the effects of ageing of upon different types of paint. An experimental development procedure is used to produce and refine a set of algorithms for an implementation of the new model, resulting in an increase in the complexity of the proposed model. The final model includes algorithms which simulate the actions of surface tension and diffusion. Details are given of the behavioural parameters and algorithms required by the model. This new model is capable of supporting the production of marks which possess a greater degree of complexity than possible with existing models. Throughout the development of the model the aim is to balance the requirements of producing a convincing visual and behavioural simulation of real paint, against the complexity of making a physically accurate simulation. The new model also provides the opportunity for new tools and techniques which are not only unsupportable with existing systems but also with traditional fine art painting methods. A selection of photographic results are included which provide support for the accuracy of the behaviour of the model.

Introduction.

Chapter 1

This thesis starts with the commonly noted observation that computer based painting systems produce images which lack sparkle when compared to those made using traditional methods and media. Their images appear like feeble or sick cousins when compared to their studio based counterparts. This thesis will discuss their symptoms, offer a diagnosis of the cause and nature of the sickness, and propose a cure. At the end of the thesis photographic evidence will show the effectiveness of the cure.

This thesis will show that the present generation of computer based painting systems are based on a simplistic model, and that this is the primary cause of their problems. This thesis will proposes a new more appropriate and elegant model for computer based painting systems. In particular it addresses the following questions:

- What is wrong with present paint systems, and why do we need a new model?
- What are the requirements of a new model?, and,
- How should one go about building one?

The thesis will describe the roots of this new model, named Wet and Sticky. It will introduce and explain concepts such as an accurate synthetic model of paint, the intelligent canvas, and the painting engine.

1.1 Requirements of the model

WET & STICKY grew out of a dissatisfaction with the image making possibilities of traditional computer based paint systems. To understand this dissatisfaction one only has to walk round almost any exhibition of paintings. If one closely examines the marks that make up the complete image it immediately becomes apparent the incredible range, variety, and complexity of these fundamental building blocks. If one now carries out a similar examination of electronic paint images one will see a much narrower range, a lesser variety, and a reduced level of complexity. These images tend to be considerably more sterile than their real world counterparts. This disparity rises out of the shallowness of the model employed by paint systems, by the lack of understanding of the process of real painting and of the behaviour of real paint.

The aim of WET & STICKY is to produce a model which will support the production of images whose marks possess the same degree of complexity and variety as those found in real paintings, and to allow for the accidental effects inherent in real painting. It is important to note that there are two yard-sticks against which the success of the model should be judged.

- The first is the final appearance in terms of complexity, variety and richness of the marks, and of the whole image.
- The second is the ease with which these marks and images may be created.

By basing the model on the artist's traditional ie. studio–based working methods, and on the behaviour and characteristics of their traditional tools and media it is hoped that both of these criteria will be met adequately.

1.2 Vested interests

Before I start I should make plain my reasons for working in this area.

My initial training was in the fine arts, and I have since taught and practiced drawing, sculpture and painting. My interest in computer painting grew out of a dissatisfaction with capabilities of the systems I used. Although these were state of the art devices, I felt they took far too shallow a view of the process of real painting. This view was supported by the fact that most fine artists I encountered who had used these systems found them too sterile and unattractive to use. My aim in proposing and developing the WET & STICKY model was to devise a system which would attract fine artists rather than repel them.

1.3 Setting the scene.

Computer based painting systems have been around for almost twenty years. Within three years of the generation of the first computer based paint image the concept of using a computer to paint an illustration had progressed from academic research to being used by the television industry. Two years later, in 1978 two dimensional computer animation, generated on a successor to the first paint system was used by NASA to illustrate the Pioneer spacecraft's encounter with Venus, (Rivlin 1986 pp.75). Today such systems are a mainstay of the television industry and with the increase in processing power and the decrease in the cost of memory chips they have percolated down from specialized turn-key systems to the home computer market. Today's top of the range systems offer millions of colours, anti-aliasing, dozens of basic tools, integration with video images and 2D and 3D animation, and many image manipulation features.

1.4 Present inadequacies.

There have been many changes in user interaction in the last twenty years: witness the changes in the way one can produce and edit documents. Nevertheless, apart from taking advantage of hardware advances,there is no fundamental difference between the first and present day colour paint system, the Shoup model, as described by Smith (1982). All paint systems use the same basic conceptual model outlined by Shoup of a frame buffer holding colour values which represent the canvas, with the user interacting directly with the frame buffer to alter and manipulate the colour values. From the most elementary home system to the most expensive turn-key system the underlying model is the same.

Unfortunately the Shoup model is fundamentally flawed as an abstraction of the real painting process.

- it reduces paint, an extremely complex and variable substance, to one single attribute, colour, represented by a simple number.
- it does not allow for any of the environmental effects that influence the behaviour, and consequently the final image, of real paint.
- it does not allow for the accidental effects common in the real paint process.
- it does not model the canvas in any more depth than as a colour.

In summary its shallow concept of paint makes it extremely difficult to reproduce the complex marks that are evident in real painting.

1.5 A new approach.

Computer based paint systems provide a simplistic model of real painting. This simplification stems from a fundamental lack of understanding of the processes and techniques of real painting. This suggests that its developer did not have any experience of real painting. With the benefit of such experience it is possible to propose a much more appropriate and accurate model. The new model should overcome many of the inadequacies of today's systems. However there would be little point in merely mimicking the real world if one did not also offer some advantages over reality. By abstracting the components, the rules governing the behaviour of paint and the processes of real painting it is possible to allow the creation of new tools which facilitate useful effects which hitherto were impossible. With the new model it is possible to directly address and interact with the attributes of paint, substrate, and the environmental conditions which influence the development of a painted image. This flexibility gives the user far more power in controlling the development of an image than is possible with existing models, or within the studio environment.

This thesis will present an argument against the Shoup model, explaining its shortcomings and reasons behind these. It will argue for the necessity of including the end users, ie. artists, in the design of paint systems. It will then go on to propose a new model, WET & STICKY. It will present the reasoning behind this model, describe its implementation, and demonstrate its advantages, which are that

- the model is based upon the attributes and behavioural chacteristics of real paint, substrate and environmental conditions such as gravity and evaporation,
- the model accommodates the accidental features encountered in the process of painting, and,
- by providing an abstracted simulation of reality, and restoring an emphasis on the process of painting, the model allows the almost automatic generation of far more complex marks than achievable with existing models.

Chapter 2 contains a review of existing computer paint systems and the associated literature. Chapter 3 describes the requirements of the new model, and discusses the WET & STICKY model. Chapter 4 is concerned with the problems associated with the implementation of the paint and canvas properties of WET & STICKY, and Chapter 5 discusses the algorithms used for the implementation of the painting engine. Chapter 6 looks at the requirements of the user interface, Chapter 7 contains ideas for further work. The final chapter contains the conclusion.

The thesis will also show that WET & STICKY provides solutions to all the criticisms of the Shoup model that have been outlined so far. Furthermore it will show that the WET & STICKY model can reproduce any of the effects that are possible with the Shoup model, and that in effect the Shoup model is a sub-set of the WET AND STICKY model.

A Review of Computer Based Painting Systems.

Chapter 2

2.1 Introduction.

This chapter shall present a review of the published work which is relevant to the topic of computer based painting systems. The sources for the literature will come from two normally disparate fields: those of Computing Science and Fine Art.

It is necessary to study the writings, and observe the working practices of artists and painters, to understand their working methods, techniques, and requirements. Since painters are to be the end users of computer based painting, and are the experts at the production of images, it would be negligent not to consider their views. Unfortunately artists are not always very good at recording their views and experiences. To artists the image is the message. Writing about technique is seen as an obstacle to their main aim, namely the production of more images. The artist is an expert communicator, but usually only in the visual sense. For this reason it is also useful to look at the work of art critics. The work of the art critic is valuable because he is one step removed from the creative process, and so can place a painting in context, to distinguish trends, and draw links between separate types of work. The role of the art critic as a documenter and describer of technique is secondary to his role as observer and critic of content, and of the aesthetic success of the image. For this reason the literature that is immediately relevant to this field of study is often buried under aesthetic considerations. Nevertheless, the aesthetic effects of certain techniques are also of importance as they can explain why these techniques were used.

A review of Computing Science as applied to graphic art is required to discover if and how it has addressed the requirements of the artist as a potential user of the computer as a tool. The review presented in this thesis will show that it is only recently that there have been any real attempts to discover the working methods and requirements of artists. This is immediately evident by the dearth of literature and research into the subject.

The review will be split into three sections. The first will give a historical perspective of the development of computer based paint systems. The second reviews the literature relating to the tools and techniques of the artist's trade. The third section will look at computer science work which has attempted to assimilate and use constructively the source material outlined in the preceding sections.

2.2 Computer based painting systems.

Before one can look at computer based painting systems it is necessary to define what is meant by the term a "computer based painting system". This will serve as a guide for categorizing different systems.

Smith (1982), Beach et al (1982), and Watters et al (1987) all give descriptions of the requirements of a paint system. From these three specific descriptions it is possible to identify elements which are common to all computer based paint systems, these are,

- Firstly their purpose; hand painting two-dimensional images displayed on some form of output device. These may be in colour, gray scale or black and white,
- Secondly the highly interactive nature of their software, often with WIMP# interfaces (Schneiderman 1987),
- Thirdly the concept of the canvas being a large piece of digital computer memory, and finally the brush as an electronic input device. It should also be noted that the user is referred to as an artist or designer.

These common features are the key elements in all contemporary computer based paint systems. Most systems offer many more features, but all rely on the same basic underlying principles. The first system that used all three elements was Superpaint developed by Shoup in 1972-73 at Xerox PARC. This system used an eight-bit frame buffer and is also noteworthy for being the first to use both the frame buffer as a electronic canvas, and a tablet and stylus as a painting input device. Smith(1982) also mentions another early system developed by Miller at Bell Labs using a three-bit frame buffer and potentiometers to change the colour of lines drawn by the user.

Shoup and Miller independently developed the first colour paint programs but the first interactive image-making system was Sutherland's(1963) prophetic work Sketchpad, developed as part of his doctoral research in the early 1960's. Sketchpad was the first interactive drawing system and the first to allow the user to communicate with the computer and directly alter the image via a light pen. However, Sketchpad does not fit our definition of computer based painting systems developed above on two accounts. His system was an object based drawing program and was developed long before the

Window Icon Menu and Pointer

advent of frame buffers and raster display devices, Sketchpad used a vector display. Indeed painting systems as we now know them only exist because of frame buffers and Shoup's concept (Rivlin 1986 pp74) of the frame buffer as a canvas. Sketchpad had no concept of a canvas. Instead it used a display list to hold the image information. The display list is simply a list of vectors. Although Sutherland's aim was to produce an interactive drawing system, not a painting system, he did however invent many of the interaction techniques that are fundamental to the drawing operations common to both kinds of systems, for example, rubber banding.

Smith's Paint program, described in his paper of 1982, used an 8-bit frame buffer and offered many more features than the original Superpaint, almost all of which are now considered to be part of the standard package of any paint system. (These are discused in detail in Appendix A). A later version, Paint3 used 24 bits per pixel, and introduced the notion of *"wet paint"*. This was achieved by using an 8 bit wetness weighting across the brush's footprint. The higher the weighting the more dominant the brush's colour is over that of the paint on the canvas.

The example given in Smith (1982) is of a brush with a random distribution of bristles with those in the centre having a higher weighting. Such a brush simulates airbrushing. The name wet paint is misleading and displays a superficial attempt and lack of understanding of the behaviour of wet paint.

A detailed chronology of the early paint systems from Shoup to Smith and beyond to the products of Richard Taylor's company, Quantel, is presented by Rivlin (1986). This work highlights the fact that with increased processing power and cheaper memory, and thus deeper frame buffers, have come more effects and more colours.

However there has been no change to the underlying model. The effects are primarily aimed at end users in the graphics and television market. Many of these effects, like Smith's *wet paint*, are simple colour manipulation algorithms masquerading under artistic names. Others provide a single effect of limited artistic potential, such as brush extrusion (Whitted 1983).

Only one paint system has broken from the mould of the canvas as frame buffer set by Superpaint. This is UltraPaint (Watters et al 1987) which is an attempt to get round the relatively poor resolution of the display device. Although reductions in the cost per bit of memory has allowed larger frame buffers, which in turn have allowed more colours to be displayed, there has been no corresponding advance in display technology. State-of-the-art displays have a resolution of 100 pixels per inch whereas print technology uses 4000 dots per inch. UltraPaint allows images of up to 32K by 32K to be created on a 1280 by 1024 display. The information for the display of the image is held as a vector representation. Using this system the whole 32K by 32K image is never stored in its entirety. Only those areas that are required for display of the present view are generated. As one zooms in, the image is always viewed at maximum screen resolution. Any changes made are stored in the vector list. As one zooms out these changes may not be visible, but can be seen if one zooms in again.

The system is very similar in principle to systems which use the PostScript page description language (Adobe 1985), but, without the advantages of PostScript's object oriented nature. Several high quality drawing programs have been developed which use PostScript, for example, Aldus Freehand (1990). The difference between programs which use PostScript and UltraPaint is that whereas UltraPaint is designed for freehand painting, the PostScript based drawing programs provide support for editable line drawing.

There is a question as to the purpose of UltraPaint. Although it allows great magnification without loss of quality, and stores the image data in a compact form, it falls down when it comes to the final use of these high definition images. If the image is to remain on the screen for final display then the only use one could see for the system would be for backdrops to animated sequences. UltraPaint would allow very large seamless backdrops to be created without the need to resort to repetition of the backdrop image. If stepless zooming were included then animation sequences that require long pans, or zooms would benefit from the system.

Willis (1984) shows how the high resolution of UltraPaint coupled with the ability to record a transcript of the users actions in creating an image can be used to produce a tool for generating high definition output for printed graphics. This appears to be the ideal role for the UltraPaint model.

Rather than just looking at previous paint systems and their models, the designer of a new fine art computer based painting system should return to basics and look at the traditional tasks and working practices of fine artists. The next section describes an informal task analysis which provides a stable foundation on which to construct a new model.

2.3 An informal task model.

2.3.1 A three point plan.

How do artists actually work? One way of answering this question is by using task analysis. Before any design work is commenced the designer of a interactive system is advised to carry out some form of task analysis. A possible approach would be that of Shepard's (1989) Hierarchical Task Analysis.

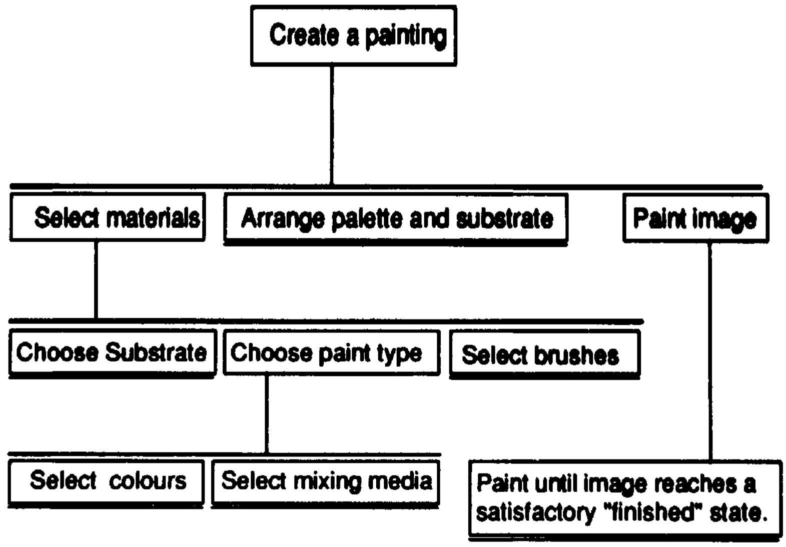

Figure 1
A hierarchical breakdown of the process of painting a picture. Operations which are underlined can be further sub-devided.

One can easily see in figure 1 how the initial subdivision of the task of painting a picture is relatively easy. The problem is that the actual process of painting the image, ie. composing the image, and the iterative process of adjusting, altering, and developing the image is much harder to define with hierarchical task analysis. As Shepard says, when it comes to analyzing complex tasks which entail considerable element of human skill, it is best to collaborate with a task expert. So as an alternative to the formal method outlined above, I propose a more informal, though thorough method of task analysis. I shall call this an informal task model.

Before designing WET & STICKY I did not carry out any orthodox task analysis, mainly because as a practicing artist for the past ten years, I am in a position where my experience, in both the area of making painting, and in speaking to, watching, and teaching others to paint, gives me suitable knowledge of the tasks to enable me to side-step traditional task analysis. Nevertheless for any other potential designer of computer based paint systems I would propose a three point plan of investigation into tools and techniques employed by artists.

- 1. become aware of the range of tools and media available to and used by painters,
- 2. review the literature written by artists, technicians and critics, and,
- 3. observe and discuss with artists the process of drawing and painting

The starting point for someone with a casual interest would be a visit to an art supply shop followed by one to an art gallery. The range of brushes, drawing implements, papers, canvases and paints is vast. Equally vast is the range of types of effects achieved with these materials. A walk round any painting exhibition, be it the work of pre-school painters or post-modern, immediately highlights the narrow range of effects possible with computer based painting.

If the prospective designer had not been put off by stage one, then stage two would be to refer to the literature. The literature can be divided into two basic areas:

- The tools and techniques employed by artists, such as Cennini (1954) and Collins et al (1987).
- The principles and concepts, such as Ashton(1972). This section would include the work of art critics such as Gombrich (1978).

Both areas provide useful insights into working methods and reasons for using one method rather than another. These two sections are discussed in greater detail below.

2.3.2 Technique.

Cennini's "Il Libro dell Arte"(1954), was written in 15th century Florence and deals primarily with the dominant painting techniques of the day, fresco and oil. It is basically a reference manual, although his advice ranges from general principles the artist should follow to specific methods for solving particular visual problems. Chapter XXVIIII sums up his spread of advice, which is entitled:

> How you should regulate your life in the interests of decorum and the condition of your hand; and in what company; and what method you should first adopt for copying a figure from high up.

Other chapters deal with the making of brushes and the uses of different shaped brushes the uses of particular colours , and methods for painting fabrics. He details methods which employ the layering of colours and how to mix colours on the substrate.

There have been many books on artistic technique since Cennini's. A modern day example is the book Techniques of Modern Artists by Collins et al (1987), which gives life-sized photographs of details from paintings accompanied by descriptions of how to achieve the effects shown. It is a useful reference as it not only shows up the deficiencies of the images produced on paint systems compared to those illustrated, it also hints at the narrow interactive capabilities of paint systems.

2.3.3 Principles and concepts.

One should not neglect the view of the practising artist despite the fact that his comments are often harder to interpret than those referring to pure skills and techniques. There also appears to be a rule of thumb that the better the artist the less likely he is to commit to paper statements about his own work. Sadly in the past there were few high-calibre artists whose conversations and statements have been recorded. However in the later two thirds of the twentieth century the artist has assumed more of a celebrity status, and consequently his views have been solicited by the media. Fashionable artists are filmed and interviewed, and are required to provide statements about almost any work they produce. All these sources, old and new can provide interesting insights into the reasoning behind a particular work, and they also can bring out ideas that may not be obvious in the final work. Such works usual refer to concepts on a higher level than those of Cennini.

My own research into this field has been influenced by those artists whose works I admire. These include Renaissance masters, the work of nineteenth century painters and sculptors such as Cezane, Rodin, Degas, and Gauguin. I have studied in detail the work of the great artists of the early twenteth century such as Picasso and Matisse, right through to those of the last thirty years. The range of styles and approaches of these artists is wide and gives a broad spread of views and attitudes to the working process.

Some of the approaches to problem solving, and methods of proceeding with a painting, may be alien to those employed by computer scientists. This is not surprising, given the differing aims and requirements of the two fields, but it is very important for the designer of a computer based painting system to be aware of these differences.

Picasso talked of the importance of the accidental in art (Ashton 1972) and of the importance of the idea as merely a starting point. He draws a parallel between walking to work and creating a work of art; he questions,

> *"On leaving your house, do you often change your route without thinking about it?...And do you not get there anyhow? And even it you don't, does it matter?"*
>
> Chipp (1968)

This wandering route is evident in the photographic and film records of Picasso at work. For example, in the series of photographs chronicling the development of Guernica taken by Dora Maar, (Richter 1965) , or the photographs of the eighteen states of the lithograph, Two Nude Women (1946), (Schiff 1986), one can see how a work may go through vast changes of emphasis and direction.

Jackson Pollock also writes of the unexplored and unpredictable path the artist follows when creating a painting. He talks of the painting having a "life of its own", and reaffirms the idea that the process is more important than the end product; he says,

> *" the instance one(a painting) is completed, the intimacy between the creation and the creator is ended. He is an outsider."*
>
> Chipp (1968 p 548)

The Dadaists often used chance as a vital element in their works. Indeed they used the idea in many different fields, from painting to poetry, from drawing to dance. The most famous being Hans Arp's *Nach dem Gesetz des Zufalls* (According to the Laws of Chance) (Richter 1965). This is a collage made by dropping a handful of coloured, roughly square, pieces of paper onto a larger piece. The fallen pieces were glued to the base piece wherever they landed, thus eliminating the artist's role in determining the composition of the image.

This desire for the serendipious influence which is an artefact of hand made image making has also been expressed by practitioners of computer based painting. On viewing some of her early computer based work, the artist Sokolove (1990) talked of feeling a need to " set human wiggliness and complexity " to disrupt the coolness of the electronic image. She views visual complexity and richness as a vital part of image making.

2.3.4 The final stage: the studio.

The final stage for our potential designer would be to talk with and observe the working methods of artists. Although this is the last stage it is also the most important. Ideally it would culminate with the designer using the artist's tools and materials. This experience would bring into sharp focus the complexities and subtlties of the tools.

By observing the artist at work, one will see painting reduced to its basic building blocks, its primary elements. Painting is fundamentally about lines and colour, and the act of making marks. It may seem over–simplistic to reduce the works of the great masters to such a basic level but in essence the act of painting, whether it be 15000 years ago in the caves of Lascaux, or in Schnabel's (1982) studio, is the same. What lies behind the marks and their effect is a consequence of their total collective interaction with each other and with the viewer and all that he brings to the painting. In essence then, all that the computer based paint system must provide is some way to make marks. Using this definition it is true that all computer based paint system's fulfil this role, but what differentiates them from real painting is the variety# and richness† of marks that can be made.

To stand in front of a blank canvas, or piece of paper, or substrate is a frightening and yet exciting experience. A virgin white piece of paper is an intimidating thing. It offers vast potential and taunts you to make your first mark. At no time subsequently will your choices be so wide and not influenced by previous events or by the collective interaction of previous marks. So the first mark may be just a throwaway gesture to break the tension or it may be a carefully placed key element in the final image. Another way to break this tension is to give the whole substrate a rough patchy tone, but this will also have the effect of reducing the contrast between it and subsequent marks. Berger (1963), gives a very elegant description of the effect of breaking the tension with the first mark. He compares a virgin piece of paper to a fish tank full of water but devoid of fish. Once a fish is introduced that is all that one looks at. The water, the paper, becomes merely the area in which it exists.

The process of painting is an iterative one. The artist may start with a plan, or idea, of how the final image will appear. Although he may well start off following this route, he cannot but help be influenced by the *process* of making the image. The accidents, the chance unpredicted composition of lines and colours, the interaction of the image with the studio environment, all have an effect on how the work progresses.

A complete sketch of the whole work may have been worked out in advance and transferred to the canvas. The sketch may serve as a very tight guide to the finished image, or may only contain a flavour of the final image or the final image may not contain any trace of the sketch. It is important to note that the initial sketch can be a valid image in its own right. The artist may terminate the work early and start afresh because the sketch would be lost or destroyed by further work. In this sense we cannot draw a parallel with a writer's rough draft. The writer's draft is rarely published and, if it survives, is only of interest to the researcher, whereas in many exhibitions the distinction between sketch and finished work is frequently blurred.

If one thinks of the marks that are made as words or phrases of a visual language then one can refer to the range and variety of the visual vocabulary used by the artist. Although one might only refer to the variety within a single image it is desirable for the artist to be able to select appropriate words and phrases from a wide-ranging vocabulary

† Using the same metaphor, variety is not enough, richness adds flavour and is an important element in enhancing the character of the work.

Throughout the painting process the painting is a living thing, not just in the poetic sense of the painting's character influencing its own development, but also at the more basic level of the physical characteristics of the paint. The paint has a life-span during which it is malleable and ranging in character from a liquid to a solid. By selecting the type of paint and medium# the artist can prolong or shorten the life-time of paints. Working methods that have been developed in the studio to solve the problems of paint drying prematurely, or conversely taking too long to dry, may not appear to have any place in an computer based paint system. One of the reasons for prolonging the time during which the paint could be worked was purely economic. The artist would mix up a particular colour and would need to use it over a period of several days or even weeks. He therefore wanted the mixture to last so that any surplus from day one could be used on day two and so on. It also avoided the problem of having to remix the same colour. This could be equated to keeping the colour values of various paints in a colour library. In a computer based paint system, however there is no danger of the paint on the palette becoming unusable due to age. Greiman (1990) makes the mistaken point that the paint never dries. This is the wrong way to look at electronic paint. It cannot dry because it was never wet.

Another, more important, reason for extending the working life of the paint, is to allow the artist to keep the paint in the image "on the move", which allows the artist the freedom to change his ideas and to mix the paint on the surface of the image rather than on the palette. For this purpose the paint's drying time need not necessarily be greatly extended, a few minutes may be sufficient.

Yet another reason is that various techniques depend on painting wet on wet. There are drawbacks to adding media to prolong the paint's life. It requires the artist to have a well-tried working method in order to be able to organize the painting process, so that he is not hindered by the wet paint.

When the artist wishes to build up many independent clean colours rapidly it is desirable for the paint to dry as quickly as possible. The most extreme example of this is *encaustic* where molten bees' wax and turpentine are mixed with oil paint. With this method the paint dries and hardens extremely quickly. The bees' wax also adds body to the paint and thus gives encaustic the rare combination of fast drying with thick gobular paint. Other fast drying paint techniques are available to produce a thin layer of paint. A more detailed

Medium is the name given to the bonding agent and/or thickener used to hold the pigment. The Oxford English Dictionary for example, defines it as *a vehicle for paint*. Examples of thickener's are linseed oil for oil-based paint and poly-vinyl-acetate for acrylic paint.

discussion of the methods and the advantages and disadvantages of altering the viscosity and drying time of paint will be given in the section on Elements and their attributes.

As the work progresses from sketch, or background tones, to more definite, perhaps finer work the artist may change the type of brushes used. As a very crude guide it could be said that the work moves from rough to smooth. This will probably happen in a series of steps. The work moves forward, later only to retreat as a result of some change of direction. This change might be caused by an outside influence, a dissatisfaction with the work, or a new idea inspired by the paint process. The exploratory nature of the working process is reflected in the way the artist typically likes to have a large selection of brushes within easy reach. One reason for this variety is to allow the artist to overcome the physical limitations of the medium. If the artist needs to apply clean colour he needs a clean brush or at least one which was last used for the same colour. The artist therefore not only needs a selection of long, short, thin fat round and square section brushes but ideally, several of each.

We should not be mistaken into thinking that only clean strokes are desirable. The streaks and smears of using a brush that has previously been used in a different colour are as valuable a part of the vocabulary of the artist as the pure, crisp strokes. Indeed, serendipious effects of this sort are often deliberately created with the artist picking up several dabs of different paint for the one brush stroke.

For the duration of the painting process different areas are being developed at different times, taken back and redeveloped. Frequently one finds that although the bulk of the image may go through several turbulent metamorphoses, patches or windows onto some of the very first marks may remain untouched. The whole image has the potential to be continually edited and re-edited. Obviously the physical properties of paint mean that this cannot be done in isolation where any element may be removed from the image and altered, and then replaced without disturbing the rest of the image.

In altering the image, the desire is sometimes to rewrite the history of the work. That is to radically review and change the previous state of the image into one more in keeping with the artist's present ideas, without leaving any trace of the previous state. On other occasions the idea is to leave a few relics of the previous state just below or poking through the surface. When parts of the image are removed or changed it may be possible to come back to a relatively clean canvas. Alternatively one might paint over the area with a colour that completely or partially obliterates what has gone before. However dependent on the properties of the paint or other substances that have been used to wipe away the substrate, or try to cover it, this will have different effects. Erasing a drawing with heavy charcoal and pencil lines will inevitably leave traces. Often these form a useful part of the subsequent composition serving as a guide and reference to both the artist and the viewer. If the area to be changed is dry then painting over the surface with an opaque paint will obscure the visual marks but the texture of previous brush strokes may not be covered. The size, direction, grouping, and flow of these marks may influence the future composition. If the area is wet then, unless the covering paint is applied in a very thick manner, the act of application will mix the top layer with the ones below. Either way, the extra thick paint or the blurred mixed paint will, at least at that stage of the development of the image, influence the artist. For many artists it is the visual historical record of a work which constitutes one of its most important elements.

2.4 Painting and computer science.

This section will look at computer painting systems which have addressed the problems of developing tools that are based on those employed by artists in the studio.

2.4.1 The designer's stand-point.

Sutherland was the first in a long line of computer scientists to develop a tool that could be used by artists. Unlike many of his successors, Sutherland, draws attention to this and sees it as a positive advantage. In his introduction he states Sutherland (1963),

> *"Had a working system* (Sketchpad) *not been developed, our thinking would have been too strongly influenced by a lifetime of drawing on paper to discover many of the useful services that the computer can provide."*

It can be argued that the development of a system can and does frequently identify many unforeseen advantages, and that a certain distance from a problem can allow a more focused view of its major elements. However it is also true that a knowledge of the process of artistic image making, be it painting, or drawing, is a positive advantage when it comes to developing an electronic (synthetic) painting or drawing system.

This is evident in the way that computer science has two very different concepts for modelling drawing and painting. In fine art there is no real definitive split between what is a drawing and what is a painting, or between how one approaches making a drawing as opposed to making a painting. Often when making a painting one starts with a drawing and as the work progresses draws into the painting. Conversely when drawing the artist often uses brushes and colours and inks and paints. However in Computer Science, there exists a very clear conceptual difference between drawing systems and painting systems.

Drawing systems

A drawing system uses an object based model and is really a tool for drawing diagrams, not for drawing in the fine art sense or in the sense of a child's drawing. To an artist a computer drawing system is a particularly alien beast. The concept of drawn objects which are re-definable in terms of colour, shape, size, and orientation, and exist in layers, is more akin to the idea of a super, if somewhat clinical, collage. Drawing systems do offer many advantages over traditional pen and paper when it comes to drawing diagrams or maps etc, but they do not in any way allow the production of fine art drawings. The use of the term drawing in their name is crediting them with much more potential than they can possible provide. Indeed it shows a very limited view of what the purpose and abilities of drawing are. This stems from the engineering and mathematical idea of drawing; for example, "I have a problem with which I need your help, let me explain it to you by *drawing* a diagram." Contrast this with the absurd image of for example, Michalangelo trying to use a drawing system to produce a drawing of the creation of Adam. Michalangelo's problems would arise from the limitation of the system which in turn stem from the definition of the requirements of the system, which in their turn stem from a limited naive view of drawing.

Painting systems

A painting system is pixel based and uses the concept of the electronic canvas. To the artist this is much more like drawing or painting. To the computer scientist the fact that it may use colour, that its tools are biased towards freehand drawing, that it not so easy to edit and is messy implies that it belongs to the world of painting. Drawing a diagram with a painting system is a slow and awkward business.

Unfortunately the role of the artistically-naive developer of computer-based paint systems has been perpetuated. This is evident in the range of tools and effects available in the contemporary commercial products, many of which have little real value other than as gimmicks such as a brush which paints a trail of musical notes. Giloth & Veeder (1985) comments on this with reference to a feature which is almost universal to all paint systems, namely the airbrush. They points out that in the studio the airbrush is used in conjunction with stencils and masks. Without these supplementary tools the electronic simulation does not begin to approach the versatility and power of the mechanical tool. Such tools are provided by the better systems but are difficult and tedious to use, whereas studio stencils are immediate and simple. Smith (1989) notes the absurd differences between what an artist may value in an image and that which a computer scientist may value. He poses the question, *"Would Picasso have been 20% better an artist if he had 120 instead of 100 colours?"*

As an artist in a Computer Science environment I appreciate that it is not quite so simple. It is true that more colours allow smoother transitions between colours, but it is also true that an image's worth, strength, and value cannot be judged by the numbers of colours used. Reffin stresses the importance of an artist's idea and the interactive capabilities of the computer rather than say, finer lines or more colours. Em (1983) supports this view and defines the key difference in approach between the artist and engineer as, that when making a work of art the emphasis is on the *process* and not on the product. Although this is a simplistic view and does not take into account the existence of process and production engineers, I believe that this is a very important point. Since the 1950's there has been a prolifereration of process–oriented art, partially fuelled by the work of the abstract impressionists, and partially by the earlier work of the Dadaist and Surrealists. Unfortunately few electronic paintbox designers have been able to see things from the artist's stand-point.

2.4.2 Modelling mark-making tools.

One can define the process of making an image as the use of mark-making tools to make an assembly of independent marks work as a whole. A mark is defined as any local disturbance to a surface (this definition allows for smudges, smears, tears, tones caused by spills, and lasting effects caused by light falling on the image, and marks caused by surface irregularities) Using these definitions let us first look at some of the mark-making tools supplied to the user of computer painting systems.

As has been previously stated, the basic tools of the artist's trade are drawing tools, painting tools and paints. One could also add to this list the various substrates used by artists for different effects. The amount of research done on these fundamental tools and the generation of synthetic versions is limited. The vast majority of paint systems offer brushes and drawing tools which are based on the same simple principle as the potato print. The brush can be thought of as the potato stamp and its shape is held as a 2D array of pixels. The canvas is a larger 2D pixel array. To make a stroke the brush image is repeatedly stamped upon the canvas along the path of the brush stroke. There are variations on this basic model, such as randomly switching on and off brush pixels as the stroke is drawn to create useful effects such as simulating drawing on a rough surface, others make use of specialised pressure sensitive input devices allied to algorithms that alter a brushes characteristics in response to changes in pressure. An example would be the brush tool in (Adobe 1990) Photoshop which maps changes in pressure from a pressure sensitive stylus to alter the darkness of a stroke. Others, such as cycling through the colour map and using a different colour for each stamp, are of little more than novelty value (although colour table cycling is a useful tool for animation)

The most notable work on an alternative to standard brushes is the work of Strassmann (1986) on *Hairy Brushes*. "Hairy Brushes" attempts to model the type of brush strokes that are used to create Japanese sumi-e paintings. In particular it models the *bokkotsu* style which is characterized by a few-well placed strokes on a light background. His reason for choosing bokkotsu was because it provides a limited range of well documented techniques to model#. Strassmann's was the first real attempt at an accurate model based on real brushes. A more recent, generalised model for calligraphic pens and brushes has been developed by Donaghue et al (1990). Strassmann's Hairy Brushes system produces very convincing, if somewhat flat brush strokes. One of the most important contributions of this work is the observation of the obvious fact that one must separate the concept of the brush stroke and the dipping or loading of the brush. I say obvious, but no other computer scientist in the field has seen this as worthy of noting, researching or modelling.

An interesting and novel approach to mimicking the complexity of brush strokes was the "drawing prism", (Green 1985). This device used real brush stokes as input to a standard paint package. The user painted with real brushes onto the surface of a transparent prism. The system did produce realistic strokes but was bulky and as Strassman (1986 pp 9) pointed out, offered no representational abstraction higher than the pixel level.

Haeberli (1990) describes a system called Painting by Numbers which allows the automatic or user driven production of "painted" images from scanned photographs. The system uses a two layer approach, with the bottom, invisible, layer being the scanned photograph which acts as the source image. By painting on the top layer the user reveals the lower image. The user can select from a variety of different brushes, each of which impart a particular style to the final image. The colour for each stroke is defined by the colour of the source image at the start point of the stroke. The system always retains an uncorrupted version of the source image, so that progressive brush strokes over an area have the effect of giving greater and greater definition to the top image. Although the system produces interesting and rich images, they are merely a skin applied over an existing image. It is not unlike the sculpting toys of the 1960's and '70's, which provided a "raw" block of plaster which the child chipped away at. Gradually this

An example is The Mustard Seed Garden, a compendium of 1000 years of sumie experience and technique first published in China in 1679.

revealed a pre-cast statue made of harder plaster that was hidden within the block. Painting by Numbers is more complex since the nature of the user's 'excavations' affect the end result, but it is essentially the same type of process.

Blesser at. al. (1988) have successfully tackled the problem of modelling charcoal. By using a tilt sensitive stylus they produced a system which appears to provide much of the subtlety and variation of marks available with real charcoal. Commercial systems are beginning to appear that can produce similar effects, (Oasis 1990).

Ware & Baxter (1989) describes the bat brush, a six-degree-of-freedom input device for modeling brushes within paint programs. The brushes work in real time and allow the user to control x, and y position, and brush size and position in the colour space. The work was done with the collaboration of artists who were excited with the advantages it offered over traditional input devices such as tablet and stylus.

2.5 Conclusion

2.5.1 A simplistic approach

As stated earlier the concept of the frame buffer serving as a form of electronic canvas is the common element to all of today's computer based paint systems. This simple yet elegant concept has survived as the only model for paint systems since their inception in 1972. However, its simplicity is also its weakness.

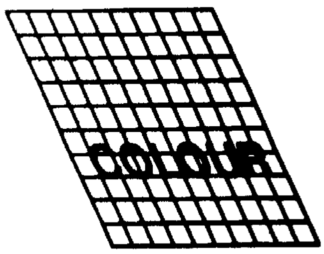

Figure 2
The Shoup Model uses the concept of the frame buffer containing colour values to represent an electronic canvas.

By reducing the concept of a painting to merely an image made up of an array of colour values the model ignores;

- the qualities of real paints
- the qualities of the substrate
- the environmental conditions that affect the paint, canvas, and painting process.

The present model has served well, and is still suitable for many applications, but its inadequacies become evident when its capabilities are compared to those of real painting. It places a priority on the end product rather than the process of painting. It does not allow the user to reproduce easily many of the subleties and complexities of real painting. This leads to an uniformity in the type of image that is produced using computer based systems. It could be argued that any medium has a particular look to it, and that computer painting is no different. This is true to an extent, but many of the traditional media used by artists allow the creation of a very wide range of types of marks, and consequently a very diverse range of image types.

The difference between computer based painting and real painting, is becoming increasingly important as more of the end users have a fine art background. To an artist used to working in oils, acrylic or water colour, the computer medium can be seen as extremely limited. This is not to say that computer based systems are not an exciting new form of expression, but that an ideal system would offer the capabilities of real painting combined with the advantages of the computer. As stated earlier, such an ideal system should look to the artists studio as its primary reference.

Immediately that one uses the artist's tools, media and working methods as a reference, the shallow and simplistic nature of the Shoup model becomes apparent. With real painting there exists a complex set of inter-relationships between the paint, the canvas and the environmental conditions. Shoup's model is only concerned with one small aspect, that of colour.

2.5.2 Colour as the final product.

All painting is partly about some form of wilful deceit performed by the artist upon the viewer. The aims and methods employed to carry out the deceit have changed through out the history of painting. In the art of the high Renaissance the primary aim of the deceit was to persuade the viewer that either, upon, within, or behind the 2D surface of the painting existed a 3D scene#. Somehow a moment in time appeared to have been captured. Perspective and highly polished realistic techniques were developed to help to convince the viewer of this. The Impressionists' aims were the same, in that they wished to record nature, but rather than striving for photographic and perspective realism they tried to capture an impression of a scene. They used different means to achieve this. They experimented with loose brush strokes and used colours in a way that created

"The function of the painter is to render with lines and colours, on a given panel or wall, the visible surface of any body, so that at a certain distance and from a certain position it appears in relief and just like the body itself." [Blunt 1]

discord. Van Gogh used lines and colours to convey what he felt about the things he painted, and what he wished others to feel (Gombrich 1978). Picasso's analytical cubist works tried to construct a reality using non–conventional means. At its simplest he used colour, line and texture to construct images which suggest the passage of time and multiple views.

Whatever the aims, in all the above cases the final method of conveying the information was a set of coloured marks upon a surface. It therefore appears perfectly valid to use Shoup's model and say that any painted image may be represented as an array of colour values. The problem with this is that it ignores the way that the final colours are arrived at, and denies the fact that, what the viewer is actually seeing is not only colour *but the colour of paint*. It therefore makes more sense to think of a painted image as an array of discrete units of paint, with each unit of paint having as part of its make up a colour value.

Figure 3
Colour should be thought of as the top layer, and only one attribute of the paint. It is the skin, of the image.

Shoup's model does not have a sufficiently deep concept of the image. The basic concept of a frame buffer serving as an electronic canvas is fine, but for it to support the working processes of the artist adequately and to model real painting more accurately, it requires the change from a frame buffer holding colour to one holding information about the canvas and its paint.

The idea is simple. A painting is just that; a surface which has had *paint* applied to it, not just a *colour* image. What one finally sees is colour, but this is merely the skin of the image. It exists as a consequence of the properties of;

- the paint,
- the canvas,
- environmental conditions, and,

• the interaction of the above.

Furthermore, not all the colours in a painting have been applied directly by the artist. Some are the result of the mixing of two or more other paints on the canvas. This mixing, indeed all the behavioural characteristics of the paint, canvas, etc. will influence the final colours. It is also important to note that not all the marks in the final image are made explicitly by the artist. Environmental effects, in combination with the paint and canvas properties, mean that the artist can step back and allow paints to mix. If he wishes he does not have to have a constant explicit control over the development of the image. It is this fuzzy control that allows the accidental which can play such an important part in the creative process.

2.5.3 The behaviour of paint.

Shoup's model is clearly lacking in one other aspect. Because it only knows about colour rather than all the contributing factors outlined above it possesses no mechanisms to model the behaviour of paint. When one is dealing with only colour there is no need for such mechanisms. Once we introduce the concept of paint as a substance possessing more attributes than simply colour it becomes necessary to provide some form of rules to govern its behaviour. An adequate model should also include the attributes of the canvas and the environmental conditions and effects. The rules should also be able to handle the way these factors influence the behaviour of the paint.

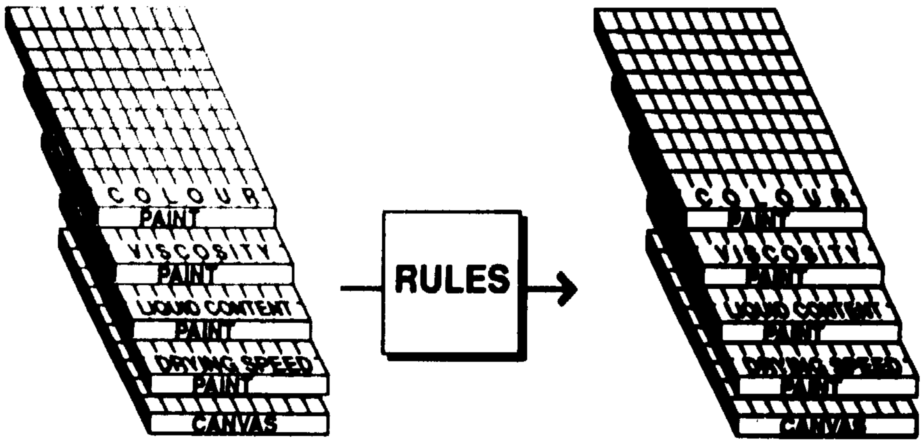

Figure 4
With reference to the paint and canvas data structure the rules transform it from one state to another.

It is clear that Shoup's simplistic model has served well over the past twenty years. Over that period there have been advances in the range and complexity of the functions available with paint systems. The benefits from the advances in hardware have also been incorporated into today's paint systems. However for the accurate modeling of the process, and subsequently the appearance of real painting, a far more complex model is required. This should take two parts;

- a data structure which attempts to completely represent all the elements that contribute to the final image.
- a set of rules which act upon and modify the contents of this structure.

Both the set of rules and the data structure should be derived from the make–up and behaviour of real paint, canvas, and environmental conditions.

The next chapter will present a description of such a model, WET & STICKY. It will also show that WET & STICKY provides solutions to all the objections levelled at the Shoup model.

Wet and Sticky

Chapter 3

This chapter will present a detailed description of the abstract WET & STICKY model. The problems encountered in implementing the model will be covered in the next chapter. The aim of this chapter is to describe all the possibilities that arise from the fundamental concepts inherent in the model, and to show how far these may be developed. The chapter will show that from a few simple components, or building blocks, a whole range of new exciting and useful possibilities can be derived.

3.1 Stage-set verses reality, facade verses simulation

Although the model is based on real world elements, the role of WET & STICKY is not that of an accurate simulation of reality in the same sense that a physicist might build an accurate simulation of fluid flow. To the artist user it is the ***apparent***, not the ***actual*** correctness of the visual and behavioural characteristics of the medium that is important. The model's responses and appearance should be consistent with the artist's expectations. We could use Hayes' (1978) concept of "naive physics" expanded to cover computation by Owen (1986), as a good way of describing the level of competence required from the model. Naive physics can be regarded as an informal, often phenomenological, precis of real laws of physics, which people develop from their encounters with the world. As long as the model responds in accordance with the artist's naive physical concepts, then how (and the exact reasons why), these responses occur, is not of importance to the artist.

To explain this, one can use the example of a stage set of a street. A stage set is a simulation in as far as it provides a thin veneer that attempts to give the appearance of reality. The viewer knows that what he is seeing is not really a street. All he sees is the facades of the buildings. They may even appear to function like real buildings, with opening doors and windows, and functioning staircases. But the structure behind and supporting the veneer is nothing like that encountered with real buildings. The similarity between the reality that the stage-set is depicting and the original need only go as far as is required to persuade and convince the viewer that, for the period or purpose of play, what he is seeing is a real street.

In the same sense, the behavioural characteristics of a model only need to appear to behave in the same way as the original. The aim is to provide just enough to trick the viewer into thinking he is observing a complete, accurate simulation, without going to the extent of actually creating one. If the trick is done well, then the model will stimulate the same responses to apparent events as to real events, (Carroll et al 1980).

However there is a big difference between the complexity of, for example the facade provided by the desk-top metaphor, (Canfield et al. 1982), popularized by the Apple Macintosh™, and the facade required for WET & STICKY. When using the Apple Macintosh it is patently obvious that one is not dealing with anything like a real desk-top. In this case the visual and behavioural accuracy of the facade is not as important as the concepts embodied in the underlying model. One can suspend disbelief in the visual appearance of the facade, because to do otherwise would entail loosing all the benefits that the metaphor has to offer.

In the case of WET & STICKY the visual and behavioural characteristics of the facade are of far greater importance. With painting, the primary channels of feedback are the appearance of the image and the way paints interact with each other. A painter with experience of traditional painting techniques and media will bring to the system knowledge and expectations about the visual and behavioural characteristics of paint. As long as WET & STICKY appears to fulfil these expectations then it can be said to be successful. If however its behaviour is erratic and not in keeping with the artist's expectations then it would be deemed a failure, and not fufilling its role as an apparent simulation of reality.

It could be said that a model which is successful in this regard could be described as a *confident* model, in that the user is confident that it will behave correctly, and consistently when compared to his frame of reference. Once a confident model has been constructed upon a simulation of reality, it is possible to make the model perform in ways which are impossible in reality but without damaging its confidence. This is similar to the concepts of literalism and magic explored by Smith (1987).

3.2 Building Blocks

How does one recreate the visual and behavioural characteristics of painting ? What are the key elements that combine and interact in the production of a painting ? If one wants to make a painting one simply requires some paint, something to apply the paint with, and a surface on which to paint. There is one thing missing from this list. Something which is not on the stock list of any art supplier but no painter can avoid using or being affected by. That is the environment. The environment affects the process of painting in four ways,

- the studio's light qualities,
- the studio's temperature,
- the studio's humidity, and
- universal environmental effects such as gravity.

We can think of paint, brushes, canvas, and environmental effects as the basic, or top level building blocks required for the modeling of painting process. To model and understand the interaction and behaviour of each of these blocks, one needs to dissect it to find the constituent parts which govern its behaviour.

To the user most of the effects of interaction may appear to occur primarily at the top level building blocks, but this is not entirely accurate. Changes at this level are merely a symptom of underlying interaction between the individual building block's attributes and between different block's attributes. As figure 5 shows, there exists two levels of interaction. On the top level are the elemental building blocks, while below them are their properties and attributes. Change appears to occur on the top level, but is in fact a result of interaction and change to the lower level.

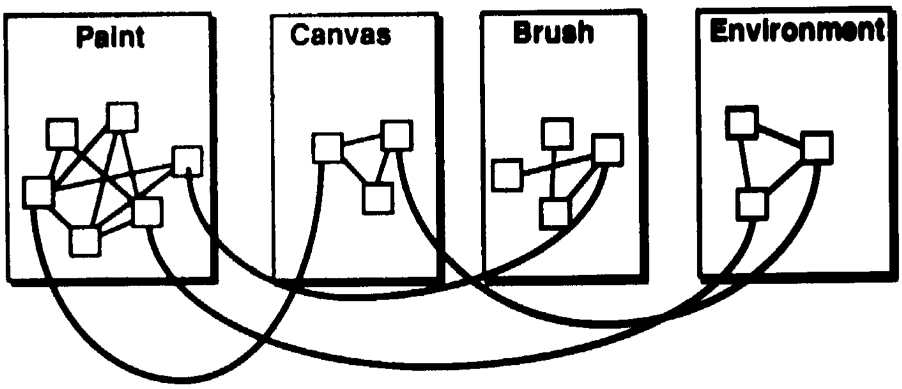

Figure 5
Interaction occurs between the attributes of single elements and amongst those of several elements.

For example, wet paint will behave differently when laid upon a saturated canvas as opposed to a dry canvas. This is due to the fact that the saturated canvas cannot absorb any more liquid and so the wet paint spreads easily, takes longer to dry, mixes readily with other paints on the surface, and does not easily adhere to the canvas. The behaviour of the brush when painting a stroke will also be different on a wet canvas. For a given quantity of paint one would get a longer stroke on a wet canvas as opposed to on a dry one. The brush would also be more inclined to glide across the surface of the canvas. All these effects which are displayed on the canvas and encountered by the brush are the result of change to just one attribute, that of the canvas's absorbency.

3.2.1 Structure

We now have our four elemental building blocks, their attributes and some form of interaction between and amongst them. The four top level building blocks are

- Paint
- Canvas or Substrate
- The Environment
- The Brush

Now the task is to identify the attributes, and how they are interrelated. These interelationships, or interactions, can be thought of as the rules which govern the behaviour of the building blocks. These rules have to be dynamic in the sense that as one attribute changes it effects change in one or more other attributes. What these changes are, why they happen and how to mimic them were the primary problems involved in developing WET & STICKY.

The next section will give an overview of the complete model, describing how the basic elements fit together.

3.3 Overview of the Model

Although the previous section outlined four basic elements, in this section I will only discuss three of them. The element that is omited is the brush. This is due to the fact that this is one area where others such as Strassmann (1986) have developed synthetic versions and for other reasons that will be explained in the subsequent text, the development of synthetic brushes was not considered within the scope of this research. However the section on further work presents a discussion of the importance and the requirements of such a brush.

The origins of the WET & STICKY model are rooted in the ideas of cellular automata, (Codd 1968) . A cellular automaton can be thought of as a large collection of interconnected finite automata. One usually refers to these finite automata as cells. In WET & STICKY these cells are arranged in the form of a 2D array of interacting array elements.

There are three parts to the model, the **Intelligent Canvas, Paint Particles**, and the **Painting Engine**. These correspond to the elemental building blocks, interaction and behavioural rules as follows;

Paint Particles	represent the paint,
the **Intelligent Canvas**	is the substrate#, and
the **Painting Engine**	is the set of rules and environmental effects that act upon and govern the behaviour and interaction of the Paint particles and the Intelligent Canvas.

The intelligent canvas is an array of cells each of which has the ability to hold paint particles. The term particle is not used to describe a particle system, Reeves (1983), but as a way of describing discrete units of paint. Paint particles may be thought of as self-contained blocks of paint. Each block contains information about its attributes, such as colour, and liquid content¤. Just like the Paint Particles, each canvas cell contains information about its attributes, such as horizontal and vertical orientation and absorbency†. The cells also know about the type and volume of paint that they hold. These cells may be thought of as open topped containers which can hold Paint Particles. The canvas is rather like a large ice-cube tray: as paint is poured into one canvas cell it will fill it up and then try to overflow into its neighbours, but unlike an ice-cube tray each cell may have a different capacity before it starts to overflow, and each cell may also have a different idea of which is the prefered direction of flow. Adjacent cells may differ in their attributes in order that the canvas may mimic different characteristics across its surface, such as sized‡ and un-sized areas. It also allows the model to mimic entirely different surfaces on the same canvas. Such an example would be a canvas which had pieces of paper or other materials stuck to its surface.

In the versions of the WET & STICKY model described within this thesis the cells that comprise the canvas map directly onto pixels on the display screen. To make a fully anti-aliased version there may be a higher ration of cells to pixels.

The artist paints onto the surface of the Intelligent Canvas with a brush in the same way that a traditional system would allow the user to paint colour onto the frame buffer.

The term intelligent is used because it also has knowledge of what is upon its surface plus its orientation

¤ These attributes will be explained in greater detail in the section on Paint Particles.

† These attributes will be explained in greater detail in the section on The Intelligent Canvas.

‡ Size is an animal based glue used to seal the surface of a substrate to reduce its absorbency.

Although in essence the brush works in the same way as a traditional simple paint system brush, in this case it now deposits paint particles onto the canvas rather than colour into a frame buffer. Any canvas cells under the path of the brush receive a number of Paint Particles. The quantity deposited is determined by the characteristics of the brush The number of particles held by a cell is known as the *volume* of paint held by the cell.

The state of a cell is determined by the state of its attributes and those of the paint it contains. By painting on the surface of the canvas the artist alters the state of the individual canvas cells. If this was the only way for the state of the canvas to change then the system would be static in between inputs from the artist and any changes to the image or state would be explicit. That is to say the artist could only change the image by applying more paint or by using a tool that affected the paint already on the surface. This would be the same as traditional systems. In a system such as GRAPHISTPAINT™ (1988) to give the effect of two paints bleeding the user has to select the paint mixing tool and use it over the appropriate area. The tool performs local averaging or in essence anti-aliasing. The footprint of the tool is very small, and it is therefore tedious to bleed paint over a large area. In reality paints will run and bleed on their own accord if their attributes permit, not because the artist uses a special tool. In the studio the painting will evolve by itself without the artist's intervention.

To mimic this independent behaviour of real paint the system has what is termed a painting engine. The engine runs as an omnipresent process continually visiting, interrogating, and updating the state of the paint held by the canvas cells.

The first purpose of the painting engine is to mimic the passage of time by ageing the paint held by the cell. Ageing the paint has the effect of reducing its liquid content, which has a direct influence upon how the paint behaves in response to environmental conditions. In effect this process is mimicking the results of evaporation. The speed of this evaporation is controlled by the paint's attributes[#] : thus allowing for different paints to dry at different rates.

Secondly, the painting engine sees which environmental conditions would affect the state of the cells which it visits. The effects it is modeling in this role can be classed as dripping and spreading effects, that is,

- the effect of, gravity acting upon the paint, and,
- a notion of sideways spreading, similar to the effect of diffusion.

[#] This will be fully discussed in the section on Paint Particles.

The painting engine looks at the state of a cell and refers to its set of rules regarding the behaviour of paint and determines if there is the potential for paint to flow from that cell to one of its neighbours. It then looks at the states of all its neighbours and decides if any of them will accept any paint. If so then it changes the cell's state and updates any of its neighbouring cells which are affected by this change.

The third way the painting engine works is in response to the application of paint by the user and the redistribution of paint between the cells caused by the dripping and spreading effects. In this role the engine is attempting to mix paint. It takes any new paint that has arrived in a cell, by what ever means, and sees if it is compatible with that already in the cell. If so, it mixes the two paints, otherwise it makes the incoming paint the surface paint.

The brush is used as the primary interaction tool. The user interface allows the user to design new brushes, to select from a range of pre-designed brushes, and to select the type of paint that will be deposited by these brushes. The model also supports the use of the existing traditional computer based painting tools, such as flood fill, editing, the drawing of regular shapes, etc. However the increased complexity and depth of the model makes these tools much more versatile and powerful. This complexity is mirrored in the complexity of the requirements of user interface.

3.3.2 Brushing on attributes

Apart from using brushes to deposit paint they can also be used to *paint* on alterations to the attributes of paint already on the canvas, and the attributes of the canvas and environment itself. It is this feature that lifts the model above merely simulating reality and onto a higher plane where new alternative realities can exist. For example it is possible to use the brush to:

- apply different gravity directions to different to areas of the canvas,
- paint on wetness, or to dry areas of the image,
- directly interact with any one or more of the attributes of all of the top level building blocks.

With such flexibility it is possible to create effects that were hitherto impossible. One can directly address the attributes of the paint, substrate and the environment. Furthermore it is also possible to *paint* on new rules so that different areas of the image behave according to differing rules.

These powerful features are only possible because the model considers paint and substrate as more than just foreground and background colours, but as complex interdependent substances, whose interaction and behaviour is governed by a set of rules based upon observations from reality. Without the depth and abstraction provided by the WET & STICKY model any alterations to a image based upon the Shoup model can only be done to the colour of that image. Whereas with WET & STICKY, alterations can be made to the underlying attributes, and changes to the colour of the image *may* occur as a result of the actions of the painting engine upon those attributes.

3.3.1 The advantages of accommodating the accidental

With traditional computer based systems the user has complete control over the development of the image. Once the user has gained experience of such a system he will be able to predict the effect of using a particular tool. WET & STICKY is a far more complex model than existing systems. However, as outlined so far, it would still be possible for an experienced user, given enough knowledge of the paints', canvases', and brushes' attributes to predict their behaviour. With a complex image made of many strokes of many different types and colours of paint such predictions would be very difficult but possible with a simple image of a few strokes. If this were the case then the model would be failing by not correctly mimicking the accidental effects inherent in real painting. In real painting the accidental is a vital part (Waddington 1969) and must be incorporated into the model.

To overcome this predictability the painting engine uses stochastically applied rules in all three of its roles. The use of a stochastic system ensures that the system is non-deterministic. For example one could apply two identical strokes of paint onto identical wet surface paint and each would produce different marks, drips and mixes. The degree by which they differ would be small but, like reality, this difference would add to the richness of the image. The only way to ensure predictability in this case would be to paint with, and onto, dry paint, in this case WET & STICKY is mimicking the Shoup model. It is this feature, that by altering the paint's, and canvas's attributes the user can alter the amount, or lack of control and predictability available to him, which is one of the model's strong points.

3.4 Conclusion

The development of a fine-art-oriented, computer-based paint system which can facilitate the production of richer and more subtle images than existing systems requires three stages. The first, the analysis of the deficiencies of existing methodologies, and the determining of how they differ from real painting was covered in chapters 1 and 2. This chapter has proposed an abstract solution WET & STICKY, which is an alternative approach. The next two chapters describes how this abstract model fares when it is forced over the hurdle of implementation. They will expand upon the requirements of the model, and detail an implementation of the model.

Paint and Canvas Properties

Chapter 4

4.1 Introduction

In this chapter I shall explain the problems encountered in implementing the paint and canvas elements of WET & STICKY. In practice, although the basic structure and ideas of the abstract proved to be sound, several insuffiencies became apparent. I will describe and elaborate on the different elements of the system, pointing out the areas that required to be developed and explaining the reasons for their development.

4.2 Overview of implementation

It was initially perceived that the way to construct a fine art based computer painting system was to build upon the best of the existing research. Therefore the first task was to re-implement Strassmann's(1986) work on hairy brushes. This was done using Smalltalk 80, (Goldberg & Robson 1983), on a Sun 4 Risc, (England 1987). Smalltalk-80 was chosen because the object oriented nature of the language provided a straightforward and appropriate implementation vehicle for the brush, bristle, and paint particle of hairy brushes. Hairy brushes also served as a useful introductionary exercise to object oriented programming. The Smalltalk environment provided a supportive structure which encouraged rapid testing and prototyping of ideas. Unfortunately the version used, Objectworks 2.5, only worked in monochrome. However this was not initially a problem as the first aim was to generate simple brush–strokes. It was found to be fairly easy to re implement Strassmann's work.

The next step was to use hairy brushes to apply a simple form of WET & STICKY paint onto an early form of intelligent canvas. Immediately wet paint was applied all the subtlties and complexity of the marks made with hairy brushes were lost. It was found that unless one was painting with dry paint, ie., mimicking the Shoup model then there was no point in using hairy brushes. The flow and movement of wet paint erased all the benefits of the mark making capabilities of hairy brushes. From that point the emphasis of the work was switched to the problems of WET & STICKY, and hairy brushes was not supported in subsequent implementations. This initial work highlighted the need for a much more complex model of brushes to work within the WET & STICKY model. A discussion of this can be found in Appendix B.

It rapidly became apparent that the supportive environment of Smalltalk was a hindrance in terms of processing and display speed. The lack of support for colour also became a distinct disadvantage. It was decided to switch to a compiled language. The subsequent versions were written in C and ran on the same hardware with an 8-bit frame store. In these later versions the primary aim was to model dripping, mixing, gravity, and spreading effects. The C implementation went through a total of 11 significant prototype versions.

4.3 Elements and their attributes

4.3.1 Paint

There is a plethora of different types of paints. Their individual properties and characteristics determine why, and the circumstances in which, they are used. Not all of these characteristics would serve a useful purpose in a synthetic model. For example there is no real requirement to model the hard–wearing properties of gloss paint. However others, such as slow drying (a property which in the studio is useful partially for economic, and partly for aesthetic reasons), should be integrated into any synthetic model. In the case of slow drying it has a direct influence upon the behaviour of the paint and consequentially the type of marks that can be made.

It is clear that those properties that are of interest relate to;

- the behavioural properties of the paints,
- the properties which the artist encounters in his interaction with the paint, and,
- those which influence the paint's interaction with other paint.

Once these have been determined one can make a generic paint which possesses the significant attributes of all paints. By altering these attributes it should be possible to produce a wide range of synthetic paints with differing characteristics. Some of these paints will behave in the same way as real paints while others will have no real world counterparts.

The generic qualities of paint

All paint consists of a coloured pigment held in suspension in a medium. The properties of the medium are used to classify the type of paint, for example watercolour, oil, and acrylic. The same pigment can be used to make different types of paint. For example raw-umber can be used as the pigment in raw-umber watercolour, oil, and acrylic paint. But this is only the first level of classification. Within paints there are many additional sub–classes. In the case of oil paint, for example, it is possible to radically alter the characteristics of the paint by altering the proportion of paint to oil medium. Some such paints are translucent while others are opaque. Some are suitable for thick impasto work, others for thin glazes. What the different media do is alter the properties of the paints. So, what are the properties that can be altered? They are,

- the paint's colour,
- liquid content,
- viscosity,
- drying rate,
- ability to mix, and ,
- transparency.

As stated earlier there are other properties which have been excluded because they have no more than a very marginal role to play in a synthetic model. These include the paint's ability to accept subsequent application of other media, its colour fastness, durability, weather resistance, and flexibility. The properties of paint pertinent to a synthetic model will now be discussed in greater detail.

Colour

To the viewer, the colour is the most important aspect of the paint in an image. To both the artist and the viewer, altering a paint's colour has an immediate effect upon their response to the paint. Traditional systems allow the artist to alter the colours he uses within a finite range of hues selected from a palette, but that is all. Although it is the final measure of an image, the processes the artist employs to arrive at that final colourful image are influenced by the paint's other properties. The methods for the display and manipulation of colours through the use of colour tables are well documented (Foley and van Dam 1990). How the WET & STICKY model handles colour mixing is discussed in the later section on the Painting Engine.

Liquid content

Perhaps the most important attribute of paint is its liquid content or "wetness". By changing the liquid content one alters;

- the way the paint will flow,
- the length of time it will take to dry out, and
- the transparency of the paint.

The first is due to the fact that altering the liquid content also alters the viscosity of the paint directly. It is obvious that the more liquid a paint contains then the more likely it is to run and drip. Even the adjective "liquid" describes something that is "flowing freely".

The more liquid a paint contains the longer it will take for that liquid to evaporate. The rate of evaporation is dependent on the surface area and the drying rate of the paint. A quantity of paint with a high liquid content spread out over a large area would dry faster than if it were spread over a smaller area. By reducing the surface area one reduces the rate of evaporation. In the WET & STICKY model the unit of surface area, the cell, is fixed thereby creating a uniform cellular unit area of evaporation across the entire canvas. Altering the drying rate alters the rate at which liquid evaporates from the cell.

The transparency of the paint is altered by the liquid content of the paint because it is related to the ratio of pigment to medium. If one increases the quantity of medium without also increasing that of pigment, then, assuming a transparent medium, the paint will become more transparent.

Viscosity

Viscosity is indirectly related to liquid content[#] . As viscosity increases so the liquid content of the paint decreases. Because of this it is only necessary to store a liquid content value only and then use it to derive the viscosity value, although the model still treats them as two separate entities. This is because, although they are directly related, they influence the behaviour of the paint in different ways.

Altering the viscosity of a paint does modify the way the paint will flow, and in this it is the same as liquid content, but viscosity does not affect the time it will take the paint to dry. In the studio it is sometimes seen to be desirable to have a viscous paint that is slow drying, and at other times it is preferable to have it fast drying[†] .

The artist has the choice of a wide range of media which will alter the behaviour of the paint. The model should not reduce this choice. In the case of drying speed, the apparent disharmony between the effects of altering the liquid content and viscosity are overcome by using a new attribute, drying rate. This exists in isolation from the other two, so altering the liquid content or the viscosity has no effect on the drying rate of the paint.

[#] The model does not take into account non Newtonian fluids such as non-drip paint which increases its flow rate when brushed.

[†] An example of a medium that is used to retard the drying process of oil paint and yet increase its viscosity, is Linseed oil. A counter example used in encaustic painting, is the addition of pure turpentine and molten bees wax.

Drying Rate

The drying rate of a paint is a measure of the rate of the loss of liquid content. This is effectively the rate of evaporation. As mentioned earlier the unit surface area is constant across the whole canvas, but the addition of different media can retard or accelerate the drying rate of a particular paint. One assumes that the paint only starts to dry once it has been laid onto the canvas. Otherwise the model would allow the paints on the palette to dry, harden and become unusable and, although this would be consistent with the metaphor, it is an undesirable feature.

There are two possible ways of mimicking different drying rates. The first is to use time-stamps. When the user makes a brush stroke, the time of the stroke is associated with each paint particle that is deposited. The drying rate would be held in the form a collection of units of time. A global drying process process would monitor the advance of time and with each "tick"of its clock, visit all wet paint, reduce its liquid content proportionally, and attend to any other propagation changes. This method can be thought of as counting down the remaining life span. An alternative would be to use an approach similar to that employed by particle systems, and to think of the paint as new-born at the moment of the stroke, and its death occurs when the number of time units past since its birth, is equal to its drying rate. The advantage of the first time-stamp method is that it allows the user to use time as a method of accessing and describing paint. For example he may wish all paint that is less than ten ticks to become dry immediately.

The problem with both methods is that they require regular continual, monitoring of the whole canvas, and, as the proportion of wet paint on the canvas increases, so would the total time required for the visits made by the drying process. Another problem is that if the wet paint is not regularly visited, drying is erratic, ie., some visits may be two or more ticks apart and therefore the liquid content would drop in jumps, and unpleasant effects can then occur. As we will see later the painting engine uses liquid content as a way of determining if an edge exists between two adjacent canvas cells. Sudden erratic changes in liquid content can trigger the engine into detecting an edge where no edge should exist.

The second method, which is the one employed in this implementation, is to think of the paint's drying rate as a measure of its *likelihood* of drying. As said earlier, the first role of the painting engine on visiting a canvas cell containing wet paint is to age or to dry the paint. In this case it does this by using the drying rate as a probabilistic measure of the likelihood of reducing the paint's liquid content by one unit. Therefore a high drying rate means a higher likelihood of the paint's liquid content being reduced at each visit. The use of a probablistic method reduces the user's control, and gives a more natural simulation of drying. Although it does not allow the use of time as a method of accessing paint it involves lower overheads and does not need to meet the regularity and precision requirements of time-stamping.

Ability to mix

Why is there a requirement to model a paint's ability to mix? At first sight it seems a characteristic of lesser importance than those already discussed. One might think that it should be grouped along with a paint's weather resistance, or one of the other real chacteristics that have no part to play in this simulation. Nevertheless it is of use to the artist. One example where it might be used would be where an artist would paint on a resist which would repel any other paint. This is like the effects produced b drawing with wax crayon and then painting over the drawing with watercolour.

Within WET & STICKY the primary purpose of defining a paint's ability to mix is as an aid to the painting engine when simulating dripping, spreading and mixing. To simulate these effects the engine requires some way of determining if paint in two adjacent cells is of the same type. If it finds that they are different it can assume that an edge or surface exists between them. This will determine the way that the paints will behave, if and how they will mix, drip, and spread. The simplest way of doing this would be to see if the two paints had the same liquid content. If they are equal then the paints could be said to be similar. They may not have any other attribute values in common, but liquid content is the most important factor governing whether paints will or will not mix. The problem with such a simple method of testing for similarity is that it detects too many edges between cells perceived as containing different paint.

The probablistic manner in which the painting engine dries paint means that two adjacent cells that originally contained identical paint may, after a short period of time, contain paints with different liquid contents. In this case the simple test would say that the two paints were dissimilar, and detect an edge. Once an edge is detected the engine tries to determine if the force of surface tension presented by the edge will be overcome by the force exerted by the paint within the cell. Obviously, the simple test would fail relatively often because it would be detecting edges within individual brush strokes, and within areas of like paint.

The ability to mix is a way of putting a range on the dissimilarity allowed between two paints before they are considered incompatible, and therefore whether an edge exists between them. The use of a range of similarity provides the user with control over spreading, mixing, and dripping that was previously impossible. A paint's ability to mix can be set so that it will mix with every other wet paint, or conversely with none, or at any position between the two extremes.

4.3.2 The Intelligent Canvas

The canvas is intelligent in so far as it has knowledge of the type of paint that is upon its surface and knowledge of its orientation. The use of the word "canvas" in the term intelligent canvas perhaps suggests too specific an image. Here the word is used to represent any substrate. The aim is to produce a generic substrate that is capable of a range of transformations wide enough to be able to mimic the characteristics of all substrates. In this way the idea is the same as that of a generic paint.

Again the question arises as to what are the attributes that are generic to all substrates? The substrate's attributes can be divided into two types: those that reflect its material properties, and those that are determined by its orientation. Examples of the first type are: colour, absorbency, texture, durability, resistance to abrasion, etc. The second type can be described by some notion of the direction of paint flow, or the direction of the force of gravity. As stated earlier the intelligent canvas also has knowledge of the strength of gravity.

I will now go on to discuss the attributes of the canvas in greater detail, starting with those that reflect the substrate's orientation.

Orientation, Gravity Direction

When an artist paints upon a rigid plane, such as a wooden board or stiff card, the direction of the force of gravity relative to the plane is constant across the whole surface. During the painting process he may tilt or rotate the substrate to influence or alter the speed and direction taken by dripping paint. Whatever changes he makes to its orientation he is bound to make them to the whole canvas. He cannot turn half of the substrate upside down, or make the force of gravity much stronger on the left-hand side of the painting. If he were to paint on a more flexible surface, say an unstretched canvas, then, the artist could to a limited extent impose some local changes to the orientation of the canvas.

The intelligent canvas's method of representing gravity's direction is sufficiently flexible to mimic both a ridged and a non-ridged substrate. Each canvas cell has its own idea of the direction of the force of gravity, or in other words, which way is down. Using this system it is possible for the artist to make *down* constant across the whole canvas, or to paint on areas which have different directions of gravity. One half of the canvas may think of the force of gravity as pulling towards the bottom edge of the canvas, and the other half towards the top. The direction of gravity is just another attribute and can be changed just as easily as any other.

Orientation, Gravity Strength

Artists often tilt the painting to alter the speed of paint dripping. For example painting manuals define 15 degrees above the horizontal as the ideal angle for painting in water colour. Although this is a good guide, often it is desirable to increase this angle to increase paint dripping and mixing. In the studio, altering the vertical angle of the canvas is the only way of effectively increasing the strength of gravity. So one can equate the strength of gravity to the tilt of the canvas. Again, with WET & STICKY each canvas cell has its own measure of the strength of the force of gravity. This provides the artist with a much more flexible way of controlling the way paints drip and flow.

The artist can target areas where he wishes to exaggerate the effects of gravity. For example, if a drip is moving down the canvas too fast, he can paint on a lower strength of gravity to slow or stop its progress. Conversely, if the paint is static in an area, he can increase the force of gravity.

Material properties, Colour

The colour of the substrate is usually thought of as merely the background onto which the artist paints. "Merely" implies that it is a simple attribute of little importance. This is not the case. The background colour can have quite an important effect on the initial development of the image.

The background may be a pure colour across the whole surface. This is typical of the standard, or most common sort of substrate, for example cartridge paper, primed canvas or sized plaster. Although a sheet of virgin, white cartridge paper is a very common surface for an artist to work from it is not always the easiest to work with. A plain, unbroken surface can be rather intimidating. To avoid this artists sometimes distress the surface with a tonal wash. The use of a non-uniform background colour gives the canvas a richer personality. Such a background is more sympathetic to the artist's initial marks.

Sometimes artists paint over old works, or photographs or pages from magazines or newspapers, all of which supply the image with an initial visual texture for the artist to work with. With WET & STICKY it is possible for the background to have either extreme, from a pure un-broken colour or any image that has been previously scanned in. The canvas colour stays with the canvas. That is, it is not affected by the paint that is applied upon it. So the colour of the background image can be thought of as being made of dry paint.

Material properties, Texture

Artists paint on substrates which possess a very wide range of surface textures. From smooth plastic, such as Formica™, to rough unsanded canvas, to the felt, fur, and broken crockery used by Schnabel (1982). Texture is a very complex subject best appreciated via the sense of touch, a sense missing from most computer interaction, (Cockshott 1989). For the viewer to see texture requires effective lighting. Painting systems exist primarily in a two dimensional, unlit world. The colours of the paint and the canvas are emitted not reflected. There is no additional light source. In a two dimensional system there is no need. How then is texture incorporated into WET & STICKY?

Although texture is a three dimensional attribute, the effect of painting onto a textured surface can be seen in two dimensions. Furthermore a WET & STICKY image contains sufficient information to synthesize the third dimension to use attributes inherent in the model as data for rendering surface attributes, such as gloss, and reflectance. Further information on this subject can be found in chapter 7.

Two dimensional visual effects

In its simplest form a texture may be thought of as a surface disturbance along the direction of the normal to the surface plane[#]. This disturbance may be positive, above or, negative, below the plane. It is a series of troughs, and ridges, high points and low-points relative to the average height of the surface. Paint applied to such a textured surface will tend to attempt to flow down from the high points, and accumulate in the low points. This leaves the high points bare, showing the background colour of the canvas at these points, while in the low points the paint forms into minute pools of paint.

The cellular nature of WET & STICKY makes it relatively easy to mimic these effects. By pre-setting the "floor" of each cell one can give the effect of a textured surface. If one uses a range from -1 to +1 to measure volume, with 0 representing the average initial volume, then any positive value represents a high point and any negative value a low point. A gradient between adjacent cells can be calculated by comparing the volume of paint held by each cell. The painting engine uses this gradient to calculate how likely it is that any paint will flow from one cell to another[†].

[#] This model cannot support complex textures ie. those that include overhangs, such as the individual hairs of a fur rug.

[†] The painting engine also takes into account the strength and direction of the force of gravity when calculating paint flow.

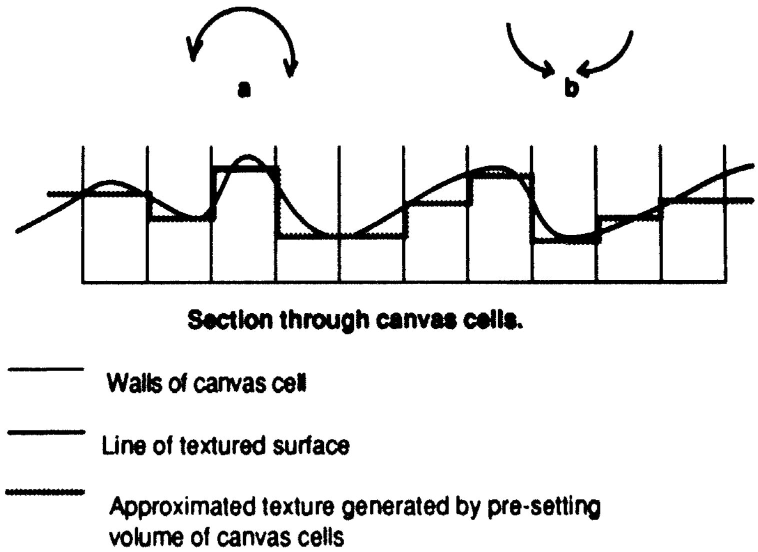

Figure 6
Paint applied to point a, will tend to flow down to its lower neighbours, mimicking the effect of vacating any high points, conversely paint will tend to accumulate in the troughs such as point **b**.

Figure 6 illustrates a two dimensional cross section through the canvas cells. Paint with a high liquid content applied to points **a** and **b** would vacate **a**, therefore letting the canvas colour show through, while paint would collect at point **b**. Figure 7 shows the same idea applied to three dimensions

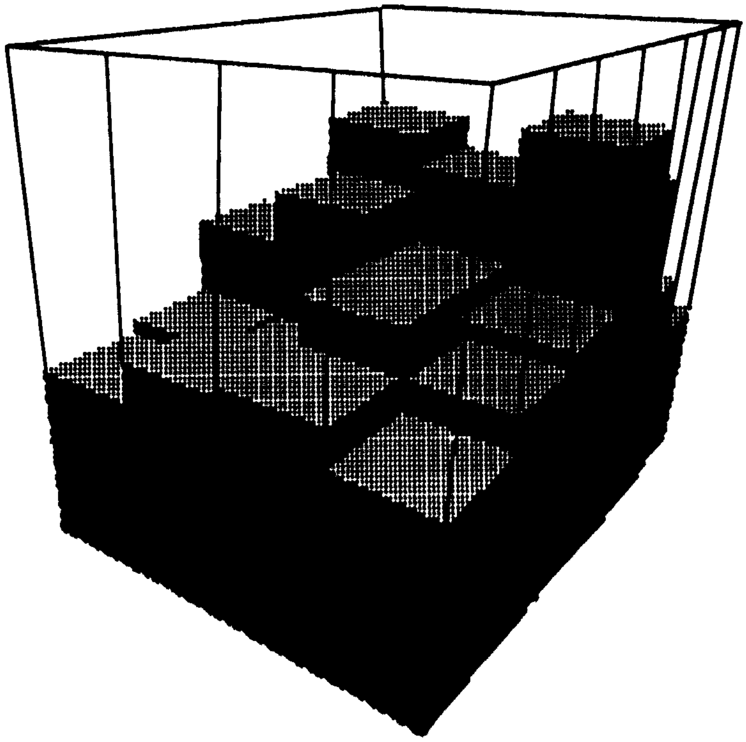

Figure 7
Simulating texture by pre-setting the canvas cell volumes. For ease of illustration all cells have a positive value.

As more paint is applied to the canvas the effect of the canvas's initial texture is reduced. Liquid paint tends to have the effect of evening out the surface height irregularities, filling in the hollows, whereas more viscous paint, remains in more ridged strokes and imparts its own texture. So although texture may initially be thought of as a canvas attribute, it is also a natural by-product of the painting process.

Material properties, Absorbency

The absorbency of a substrate has the effect of reducing the liquid content of any paint applied to its surface. The absorbency of a substrate usually only has an effect on the first application of paint. Unless the substrate is very absorbent (un-sized plaster would be a good example) then the first application will contain enough liquid to water-log the substrate, and subsequent application will not be affected by the substrate's absorbency. Absorbency can also be thought of as a way for a cell to retain a minimum quantity of paint. If we return to the ice-cube-tray metaphor, then absorbency can be equated to the height of the walls surrounding each cell. However, there are also more complex aspects to absorbency.

Dependent upon its properties, once the first application of paint has dried, and the substrate has dried out, absorbency may effect subsequent applications. For example, in the case of water colour on heavy weight paper, the dried water colour presents a permeable layer through which the paper can absorb the liquid from subsequent applications of paint, whereas with oil paint upon canvas the dried paint forms a non-permeable layer which renders the surface non absorbent.

To model this more complicated aspect of absorbency the model would have to have some concept of a paint's permeability. The model as described does not support complex absorbency, as it is believed that to do so would be of marginal benefit. However the effect of complex absorbency can be obtained by the user painting on absorbency. At any time during the painting process the user may paint on any attribute at any level. The user may for example, paint a circle of very high absorbency and a square of nil absorbency. This has the effect of re-setting the chosen attribute of all of the canvas cells under the brush's path.

4.4 Conclusion

This chapter has described the attributes of the paint particles and substrate and discussed the problems associated with their modeling. It has highlighted liquid content as one of the most important attributes governing the behaviour and interaction of paints. An important point it has introduced is the concept of employing probablistic methods to provide naturalistic modeling, in this case of a paint's drying rate. Although no mention is made of specific numeric value ranges for any of the canvas or paint attributes, experience from implementing several versions has shown that they should lie between 1 and 100. Only experimentation with a specific implementation provides information as to the exact working range, or calibration of any attribute.

Painting Engine

Chapter 5

The painting engine is the heart of the model. It is the element that lifts the model out of the mould cast by the Shoup model. All the elements described so far are merely there to fuel the painting engine. Without them it cannot operate, and without the engine there is no point in modeling paint and canvas attributes.

The engine continually executes a set of rules. These rules are based on observations of the real behaviour of paints and substrates. The aim with any implementation of the painting engine should be to try to construct as compact and as simple a set of rules as is possible, yet one which gives a convincing simulation of reality. The set of rules should be compact because they are repeatedly applied to every canvas cell or at least every canvas cell containing wet paint. The smaller and less time–consuming this process can be made then the faster the system as a whole will perform.

An important point about the painting engine is that it always works from the perspective of the host cell. The host cell is the cell which the painting engine is presently executing its rules upon. The painting engine looks at a cell and determines how the cell will interact with its neighbours, rather than how they will interact with it. The host cell may donate paint to its neighbours, and not vice-versa. If the engine applied rules to the host cell which determined and implemented the behaviour of itself *and* determined and implemented the behaviour its neighbours, then the engine could drive a cell into a state of flux. For example, it might decide that the host has surplus paint which it could donate to a neighbour, this could alter the state of the neighbour sufficiently so that it had surplus paint, and donate it back to the host and so on. As long as paint may only flow one way, from the host cell to its neighbours, the engine avoids oscillating effects.

This section describes the different aspects of a basic painting engine's task and the implementation of a painting engine. Figure 8 shows the sequence of the various parts of the painting engine's cycle.

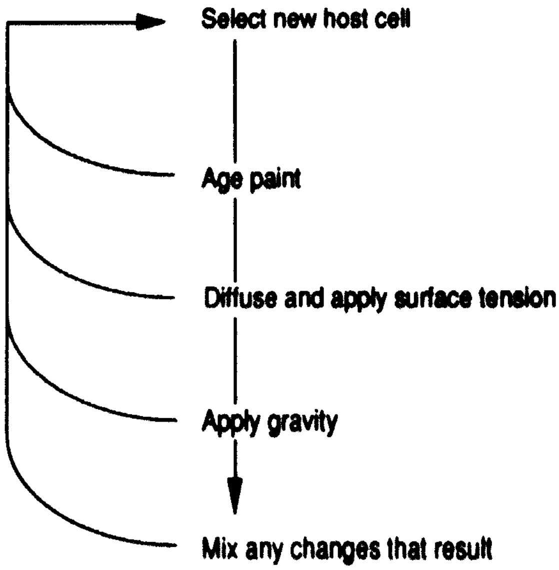

Figure 8
One cycle of painting engine. The engine may only perform a partial cycle dependant on the effects of ageing, diffusion, or gravity.

5.1 Selecting host cell : visiting schedule

In traditional paint systems time stands still: there is no concept of the passage of time influencing the behaviour of the paints and appearance of the image. The role of the painting engine is to model the environmental effects upon the paint and canvas through the passage of time.

In the studio, time passes at the same rate at every point on the canvas. A real painting exists in a parallel processing environment where no part of the image is ageing at a faster or slower rate. With a synthetic model one can attempt to mimic a universal unit rate of time passage in two ways.

- sequential scanning, and
- parallel processing

Each approach has its advantages and disadvantages, I shall now discuss both and argue for the use of a type of synthetic parallel processing.

5.1.1 Scanning

The first method is to scan the entire canvas from left to right and from top to bottom. Visiting each cell in turn, performing all the evolutionary# processes upon that cell, and then moving on to the next. There are two problems with this method.

Firstly, the length of time it takes for the painting engine to carry out its tasks at any particular cell will be determined by the state of the cell's characteristics and that of its neighbours'. When it comes to updating the display either this is done on a cell by cell basis, with the display update proceeding erratically across the screen, or, alternatively, as a complete refresh once all the canvas cells have been visited. In this case the image would appear to evolve in jumps. The effect is similar to watching dancers illuminated by a strobe light.

Secondly, it was found that the first implementation which made use of a scan line method resulted in an unrealistic movement of paint. With a canvas with a uniform direction of gravity from the top to the bottom, as the painting engine visited all the cells on the top row, it attempted to shift any surplus paint down to the row below. Row by row the effect was the same, with a wave of surplus paint preceding the downward moving scan line. With this type of scanning method the scan line acts like a windscreen wiper, sweeping surplus paint ahead of it. The result is that in just one scan, the bulk of the surplus paint from the whole image will have collected in the bottom row of canvas cells. Although it is harder to predict the effect on canvasses with non–uniform gravity direction, the effect on what can be considered the standard type of canvas, ie. with uniform gravity direction, is unacceptable.

An evolution, is one cycle of the painting engine, that is a complete, visit by the painting engine, an attempt by it to change the state of the paint held by the canvas cell, and an update of the cell and any of its neighbour effected by a change in the state of the cell.

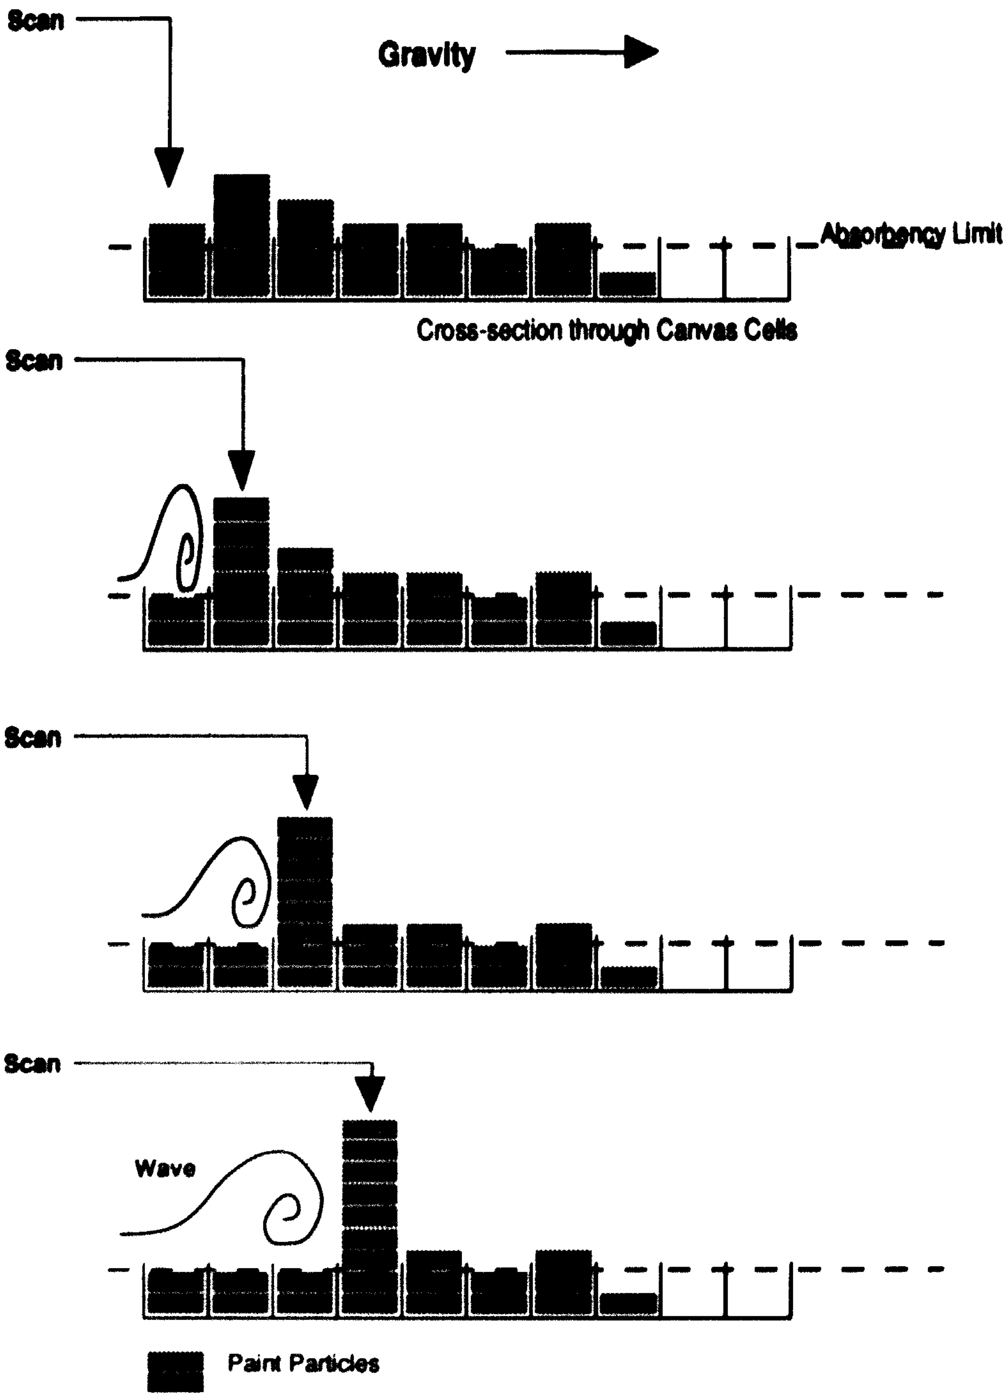

Figure 9
Illustrates the "wiper" effect which results from the painting engine using a baustrophedon visiting schedule

Thirdly, there is the problem that sequential scanning can induce the particles of paint to slosh between adjacent cells. For example, if upon visiting cell(x,y)# the engine passes its surplus paint to cell(x+1,y), then on the next cycle of the engine it visits (x+1,y) it could decide to return the same surplus paint. However because this transfer takes place over two cycles, the ageing process would mean that paint's liquid content may have altered. Therefore although the cells' volumes of paint would be unaffected by the interaction of the engine, the characteristics of the paint they contained would have changed.

5.1.2 Parallel processing

Parallel processing decreases the length of time taken for the painting engine to perform a visit to every canvas cell. Rather than just one processor carrying out the tasks of the painting engine there are two or more each with its own set of cells and instance of the painting engine. Taken to its logical conclusion, each canvas cell would have its own processor.

But even with a processor per cell, time would not pass in a uniform and constant manner. Certain cells would still require the painting engine to perform more calculations than others. A cell containing no paint could not possibly evolve, and therefore its painting engine would be idle, whereas a cell with a lot of wet paint would require its painting engine to perform many calculations per evolution. Unless restrained, certain cells would evolve at a faster rate than others. Whether this would be a problem or not is uncertain. Certainly it would result in uneven drying of similar paint, which would result in the problems of the painting engine detecting paint boundaries within strokes. One way to overcome this possible problem would be to define a unit of time that was equal to the length of time required for the most complex evolution. This could then be used as the "tick" of the clock, and with each tick all the processors would display the current state of their canvas cells, and start their painting engines. Apart from the potential speed advantage of a parallel WET & STICKY, a parallel approach removes the wiper effect inherent with the scan method.

5.1.3 Quasi–parallel processing

In an attempt to mimic parallel processing the model employs a random visit system. Instead of scanning the canvas, the painting engine selects cells at random. This means that in WET & STICKY time is not measured in units of constant length, rather in terms of

---

x being horizontal and y being vertical.

visits to each canvas cell. The engine visits a cell, evolves through one cycle, updates any resultant changes to itself and to the states of any of the host's neighbours, and then displays any changes to the colours of the paint in the affected cells. This method avoids the wiper effect and gives the impression that the canvas is evolving at the same time all over its surface. Employing a random visit also means that on average all the cells will be visited with equal regularity. Of course this is only true assuming a uniformly distributed random number generator. However it does mean that at any one moment adjacent cells may have been visited with different regularity. Again this is the same problem that occured with earlier implementations. The problems that arise from these localized differences and how the model solves them, are discussed in the sections, Ability to Mix, and Surface Tension.

5.2 Ageing

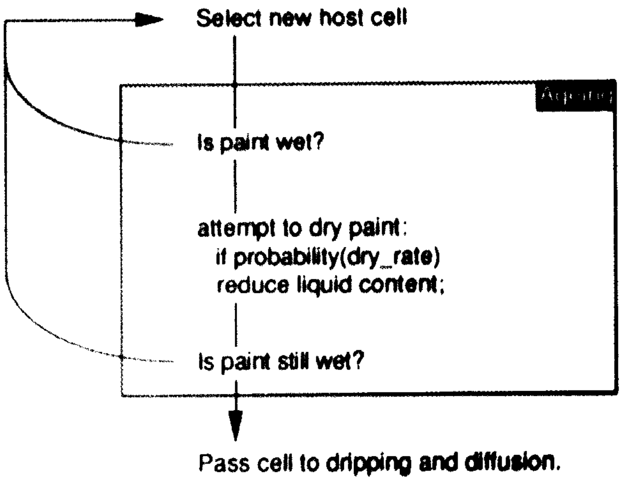

Figure 10
Illustrates the ageing part of painting engine's cycle upon a canvas cell.

As figure 10 shows, the first task of the painting engine on visiting a cell is to determine if it contains any wet paint. If the cell has no paint or only dry paint then the engine moves on to the next randomly chosen cell. If, however, there is wet paint present then the engine attempts to age the paint. As mentioned earlier, the success or failure of this action is determined by the paint's drying rate. A high drying rate means that there is a higher chance that the engine will be able to dry the paint, and conversely, a low rate implies a lower probability. If the engine is successful it then reduces the paint's liquid content. This mimics the major effect of the passage of time upon paint, that is, the gradual drying of the paint. As the liquid content is reduced so the viscosity of the paint is increased.

In reality the passage of time also affects paint in other ways. For example, as encaustic paint dries out it also tends to become more brittle, increasing the likelihood that subsequent brush-strokes will cause the initial layer of paint to snap off from the substrate. This can cause whole patches of paint to fall from the canvas. This effect is considered undesirable and is not modelled, as it would require the model to represent and record brush-strokes as discrete and repeatable units. It also presents the problem of how to represent in a two dimensional view the sudden unexplained appearance of areas of bare canvas.

Another common effect of drying out is the formation of a skin over the surface of the paint. Although this is not modelled directly, any increase in the viscosity of a paint also increases the force of surface tension resisting any movement of the paint. This effectively makes the paint behave in the same way as if a skin was forming.

If, after this drying process, the paint still has some liquid content then the painting engine moves on to the next part of its cycle, the movement of paint under the influence of gravity, and redistribution of paint by diffusion. Figure 8 (on page 57) shows a complete cycle of the painting engine. If ageing has the effect of completely drying the paint then the cycle ends at that point, or, if all the effects of dripping and diffusion cause paint to flow into cells containing only dry paint, or into cells that are empty, then no mixing need be performed, and the cycle is terminated prematurely. Depending on the initial state of the cell and its neighbours, the engine may only perform a partial cycle.

5.3 Dripping and diffusion

To mimic dripping convincingly the model has to take into account not only the effects of the force of gravity upon the paint, but also the effects of the surface tension of the paint, and the diffusion of paint. It was found that if one tries to treat gravity separately from surface tension and diffusion the results are unrealistic and simplistic. In the early versions of the second implementation, the system only modelled gravity. Without any model of surface tension the paint just fell straight down in the direction of gravity, and its progress was only halted when the paint dried up. The next version introduced a simple model of surface tension which attempted to make the paint flow within the boundaries of an area of similar paint, i.e. with the same or similar liquid content. It did this by ordering the directions that surplus paint could flow. Priority was given to the cell directly below, but if that contained dissimilar paint, or was empty, then it tried to flow to the lower left or lower right cell. Again if their contents were dissimilar, the surplus paint would attempt to flow horizontally, left or right. Only if none of its neighbours (discounting those above the host cell) contained similar paint, would the engine pass the surplus to the cell bellow. This model of gravity and surface tension resulted in drips only forming at the very bottom tips of strokes and at isolated downward pointing peninsulas along the side of the stroke. Because the model prefered to move surplus paint to the cell immediately below, the drips that formed were always only one pixel wide. It was found that a much more complex and interlocking model of gravity direction, strength, surface tension, and diffusion was required.

These four elements are so tightly interconnected that it is hard to know which is of most importance, and therefore merits being described first. I shall therefore describe them in the order that the painting engine deals with them.

5.3.1 Diffusion

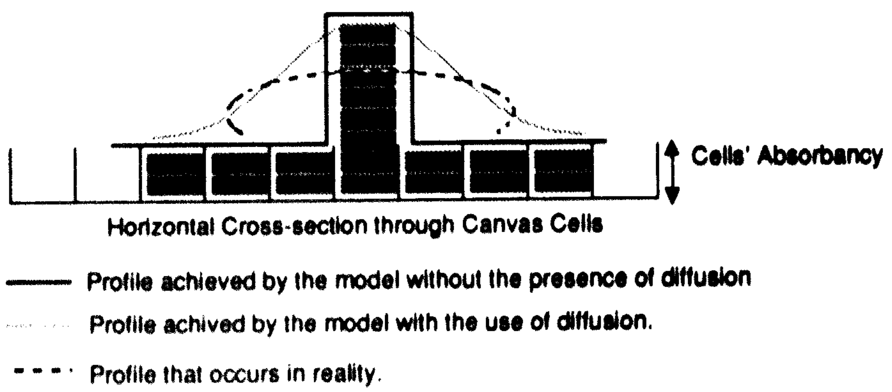

Figure 11
Illustrates why the system has to model diffusion.

Without a model of diffusion it is possible for paint particles to form into isolated tall stacks. These stacks can be formed by the action of the user painting a stroke onto the substrate, or by the cumulative effects of gravity. Without diffusion the only way to lower these stacks to a more realistic level, ie. to remove their surplus paint, is for gravity to push wet paint particles from the stack into the cell below. In such a simple model there exists no force to disperse paint particles to the neighbouring side cells.

If one employs a model which includes a simple notion of surface tension along with gravity strength and direction, then the formation of paint particle stacks is limited to the interior, and to straight, lower, edges of patches of similar paint. This is because, at any angled edges, the effect of surface tension is to attempt to spread surplus paint sideways as well as downwards.

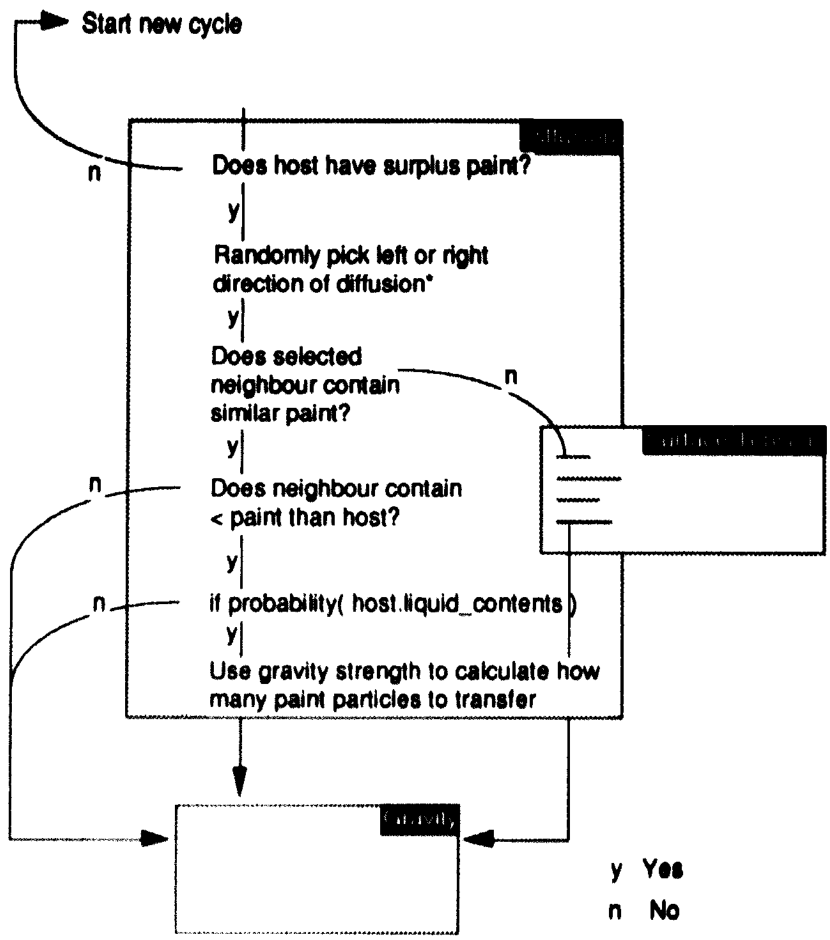

* The engine checks for invalid cells, ie. those that lie off the edge of the canvas

Figure 12
The interrogation of a cell and its neighbours as the painting engine calculates diffusion. The probability function uses the host cell's paint's liquid content to calculate if the paint has the potential to flow.

Figure 12 illustrates how the present implementation models diffusion. It is necessary to pick at random a left or right direction of diffusion, otherwise paint tends to creep slowly in one direction or the other. One determines if two cells contain similar paint by comparing their paint's liquid contents and the host cell's ability to mix factor. This provides a range within which any differences in liquid content will be accepted. Although a negative response to the question "Does neighbour contain less paint than host?", terminates the diffusion element of the painting engine, it is still possible for gravity to have an effect upon the host cell. The reason for this apparent abnormality is explained in the next section.

5.3.2 Gravity

Gravity has the effect of always trying to move any surplus paint to the neighbouring cell that is defined by the host cell's gravity direction attribute, ie. the cell that is downhill relative to the host. The WET & STICKY attribute *gravity strength* is a measure of the force exerted to move paint.

A simple approach to modelling gravity would be to assume that paint would only have the potential to flow from one cell to another if the lower cell contained less paint. This is also assuming that both cells contain similar paints. (We can make this assumption, because, if their contents were different, the engine's surface tension procedure would detect an edge and the resistant force of surface tension would come into effect.) This method would work if all cells were influenced by gravity at the same time. However, since the model uses only one processor, one can only model the effects of gravity on one cell at a time, and this method will not work. The result of using such a simple approach is the wiper effect discussed in the section on scanning.

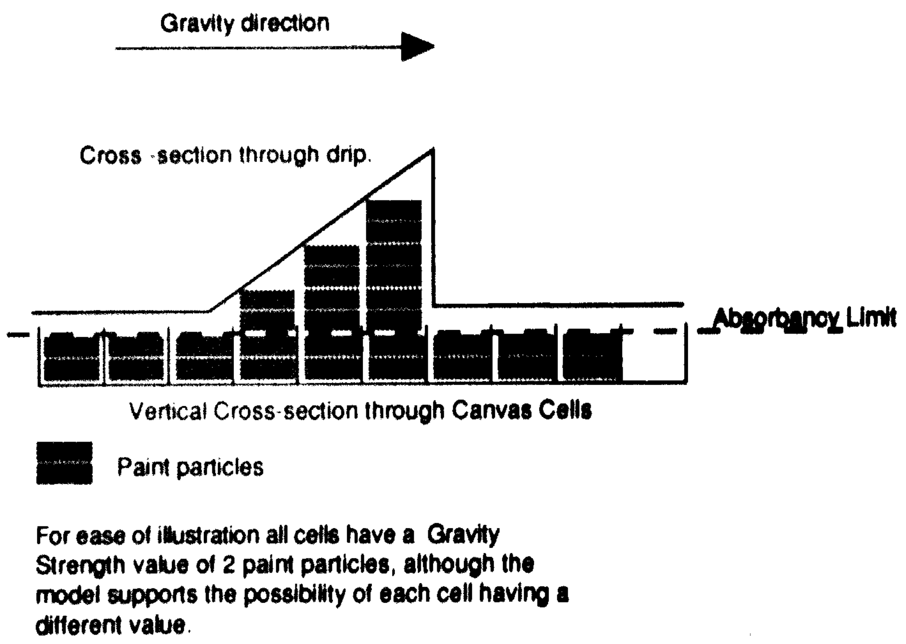

Figure 13
Illustrates how gravity strength is used to control the trailing slope of drips.

The natural profile of a drip has a leading edge which is taller than its trailing edge. The simple approach would produce drips which had a profile which would tend to have a tall trailing edge. WET & STICKY uses gravity strength as a way of regulating the gradient of the rear slope of any drips that form. Figure 13 illustrates how this is achieved. The figure illustrates the case where all cells have gravity of the same direction and strength. Gravity strength is used to determine whether the host cell has enough energy/force to move surplus paint to the cell that is downhill from it. The maximum quantity of paint that can be transfered is half of any surplus, ie.

$$\frac{1}{2}(\text{ host's contents - host's absorbency })$$

A host cell has the potential to move a surplus paint particle to its downhill neighbour if,

$$\text{downhill cell volume} < (\text{ host's volume + host's gravity strength })$$

The final factor determining whether paint will flow is the viscosity of the host cell's paint. This is used as measure of the likelihood of paint, with the ***potential*** to flow, actually managing to flow.

This method of simulating gravity ensures that on a canvas with uniform gravity, the leading edge of a drip will always tend to evolve into a state where it contains more paint than any other part of the drip, and that the trailing edge will assume a gradual slope which is regulated by the strength of gravity. This is not a totally correct simulation. In reality, a drip's leading edge may form a bulbous overhang. To achieve this effect would require the model to support the concept of a canvas cell possessing some record of contents of the space above its surface. Like the opposite of mining rights, each cell would have control over any paint that entered its air space. But if one uses the guideline that the model only has to appear to be correct, then there is no need to increase the model's complexity to include the concept of "air-rights".

5.3.3 Surface tension

Surface tension has the effect of resisting the movement of any paint at an edge. As stated earlier an edge occurs where there exists a difference in the liquidity of the paint held in any two adjacent cells. This obviously also applies when one of the two cells does not contain any paint. Another effect of surface tension is to tend to force the paint to assume a shape which has the shortest possible boundary. Consequently one only ever sees drips with a curved profile, one does not see pixelated, jagged, rectangular, or triangular drips. A system based on the model as described so far produces drips which form into a triangular profile. An early version of WET & STICKY also ignored diffusion and the resultant drips, which were one pixel wide collected at the bottom of thepaint stroke.

Using high resolution screens and with the subtle modulations in colour that occur due to the simulation of mixing, the model can produce drips that appear to be curved. These modulations occur due to the fact that the initial flow of paint from one area of paint into another allows the transfer of only a few paint particles. Such small quantities cause very subtle mixing, giving a kind of localised anti–aliasing. To accomplish curved drips the model employs a single concept of surface tension that embodies both the effect of resisting paint flow and curving the profile of drips.

Once the engine has detected an edge it has to determine what is the profile of that edge. With a square pixel based display there are only four distinct profiles.

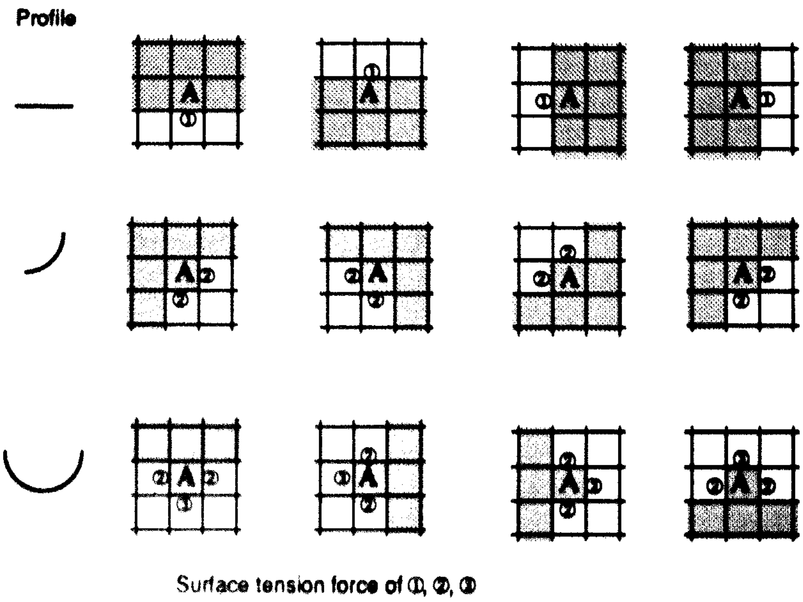

Figure 14
Three basic edge profiles of a cell A, in all four orientations, and the force of surface tension against each edge. The fourth profile, that of an isolated cell of paint, is not illustrated.

- (first case), where the cell is in line with its immediate neighbours and can be said to be at a straight edge,
- (second case), the case where the cell is on a staggered edge,
- (third case), when the cell is protruding from its surrounding cells,
- (fourth case), when the cell exists in isolation from similar paint and is, in effect, a special version of the third case.

Each of the first three cases crudely approximate a straight edge, a convex side, and the convex tip of a drip. Concave edges are derived from a combination of the first and the second or third case. A single cell cannot model a concave edge. This requires the profile of two adjacent cells to form a shape where their surface tension forces co-operate to mimic the effect of a concave profile. One edge will have a surface tension force of two, while the other's would be one#. The number of paint particles held within a cell is a measure of the force that is exerted by the paint against any edge. A concave surface is weak and offers little resistance to the force exerted by a cell's volume of paint. Whatever the orientation of the concave profile, either the effects of gravity, or diffusion, or weak surface tension, will result in the surrounding paint having a high probability of flowing into the vacant cell.

To calculate force of surface tension acting against a particular side is a two stage process.

> First the model, checks to see if that side is abutting onto dissimilar paint. If true then it proceeds to the second stage, if false, then that side has a surface tension force of zero.
>
> Second the model counts the number of adjacent sides that are abutting onto cells that contain dissimilar paint.

For example, if we look at the first row of figure 14, the case of the straight edge, we can see that only one side of cell A abuts onto a cell containing dissimilar paint. This gives that side of the cell a low surface tension figure of one, and all other sides a figure of zero. A figure of zero implies no surface tension, and therefore paint may easily flow from the host cell to its neighbour. As the surface tension figure increases, so does the resistance offered by an edge. The straight edge is a weak structure offering little resistance, and a figure of one implies that only a small force is required to break through such an edge. However, in one cycle of the painting engine, diffusion only acts on one vertical side. Therefore in the 3rd and 4th examples on row one there will be a 50% chance that paint diffusion will cause paint to flow into the similar paint rather than attempting to break into the dis-similar paint. This means that on the vertical sides of drips, paint spreads sideways slowly, if at all, even though the resistant force of surface tension is low. This effect is enhanced by the fact that gravity moves paint to the bottom of a drip, resulting in the vertical sides of drips having a lower volume, and therefore they are unable to exert as great a force to overcome surface tension.

In the second row, two of A's sides abut disimilar cells. Both of these sides have a surface tension figure of two. In the third row, surface tension applies to three sides. If we look at the first example in the row, we see that the lower side will have a surface tension figure of three, while the left and right sides will have a figure of two.

These values are just relative, not absolute values.

Once the painting engine has determined that a cell has surplus paint, and has the ability to flow in a given direction, be it under the influence of diffusion or gravity, it has to determine if it also has the potential to overcome the force of surface tension.

The potential to flow is calculated by the host's liquid content multiplied by the host's volume which itself is modified by the surface tension figure. It is important to note that volume is used to calculate potential to flow, ie the greater the volume then the greater the potential.

potential = host's liquid content * (host's volume - (selected side's surface tension))

This potential figure is then used as a measure of the probability threshold that paint will flow. Given that a cell's paint has the potential to flow, and the probability function has returned a positive response then before paint will flow the engine must again check that the cell contains more than its absorbency limit. This is necessary because the process of diffusion may have already denuded the paint's volume to below the absorbency level

In summary, the greater the volume and the more liquid the paint then the more likely that the force exerted by the paint will overcome surface tension. This is then combined with the following elements

- in any one cycle of the painting engine, diffusion only has a 50% chance of acting on a vertical edge, this is due to the fact that within one cycle the painting engine only acts upon one vertical side
- the general effect of gravity to make paint collect at the bottom of drips, and
- the weighting given by surface tension,

Paint will then tend to form into realistic, curved, attractive–looking drips. This system requires that for the formation of drips a cell needs the co-operation of its neighbours. For example a cell that breaks through a horizontal straight edge, alters its profile to that of a solitary protruding cell. With this profile it finds that surface tension hinders its downward vertical motion and favours its sideways motion, thus tending to broaden the tip of the drip, altering the surface tension of its neighbours, and consequently their behaviour. Therefore without even directly interacting with a neighbouring cell, ie. donating paint, the behaviour of one cell influences that of its surrounding cells.

A final point to note is that the model does not take into account the effect of surface tension upon the top side of a cell. The model assumes that the effect of gravity would negate any upward movement of paint, and therefore there is no need to calculate a resistive surface tension force.

5.4 Mixing

When new paint is introduced into a cell, be it by brush, diffusion, or gravity, the model mixes paint if the paints have similar liquidity. The model mixes all the attributes of the paints proportionally according to the relative volumes each of the two paints will contribute to the new paint.

In the case where the two paints are dissimilar the new paint sits on top of the old. The old paint underneath takes on the properties of the new top layer. This is obviously not realistic. A more realistic system would be to keep the paint separate in layers, and have the painting engine act upon each layer independently, although with this system the engine would have to model the mixing that might occur in the future between vertical layers of paint. The current model only supports canvas cells possessing one type of paint. The layering that results from combining dissimilar paints could be incorporated into the model, but the increased complexity and decrease in speed this would bring, and the vast increase in memory that would be required# , would not be offset by the visual advantages derived from modeling accurately this aspect of mixing. Again one comes back to the idea that the model should appear to be correct without actually going to the extent of being correct.

Colour mixing is also handled proportionally, but where most paint systems employed an additive mixing system, Wet & Sticky uses a subtractive one. In reality pigment mixing is subtractive, ie. red, blue, and yellow make a dark browny black, rather than white as in an additive system. The model uses the HLS colour system which is one of the most accurate in its mimicking of subtractive mixing, (Foley and van Dam 1990). In the present implementation colour values are thought of as points in a three dimensional colour space. A vector drawn between any two points represents all the possible mixes between the two colours. The relative numbers of paint particles of each colour in the mix are used to calculate the distance along this vector. However, due to hardware limitations the current implementation only uses a limited palette. The current implementation works on a Sun4 Risc work station with an 8 bit frame buffer and colour monitor. Rather than use the whole spectrum, and have a limited range of intermediate tones for mixing between for example red and blue, it was decided to employ a limited palette. This allows for finer gradations between colours.

For each additional layer of paint a canvas cell would have to store an additional record of that new paint's volume,and all its attributes.

Although our hardware limitations do not allow the present version to model transparency, only small changes would be required to do so. Either one could add an alpha channel to the colour value of paint particles or one could use the more complex physical based model of paint mixing proposed by Oddy & Willis (1991). Although this second approach is appropriate to WET & STICKY, and could be incorporated into the model, to do so would require some additions to the basic model. The physical colour model is primarily concerned with static colour mixing while WET & STICKY is concerned with dynamic mixing, and has no notion of the way in which the thickness of a layer of paint would effect its transparency. In general the transparency of paints with a high liquid content tends to decrease as the thickness of the paint layer increases. This is particularly evident in the formation of drips of water colour paint. Whereas with viscous paints increasing the paint thickness has little effect on transparency.

These effects could be modeled if one used the paint's liquid content to modify the proportion of pigment relative to medium, that was passed by by the painting engine under the influence of gravity. For example, above a certain liquid content threshold any paint particles passed would contain a higher proportion of pigment to medium than that held by the host cell. This action would also have the effect of reducing the amount of pigment remaining in the host cell. The end result would be that the host cell would become slightly more transparent, because it contained less paint, and the receiving cell would hold more paint and become more opaque. Alternatively, both the volume of paint held by a cell and the liquid content of the paint could produce a transparency value.

A similar effect could be implemented using liquid content to modify the alpha value of donated paint.

5.5 Conclusion

This chapter has highlighted the advantages of using a random scanning method, the use of a probalistic method of modeling paint drying, and introduced algorithms for diffusion, gravity, and surface tension. The integration of all these elements produces a realistic simulation of dripping. Perhaps the most important element is surface tension. The use of the simple yet elegant idea of modeling on a micro level the forces of surface tension, when applied universally across an area of similar paint induces a realistic response on the macro level.

User Interface

Chapter 6

6.1 Introduction

Any computer based painting system requires an interface that exists on several levels. One can separate these as follows;

- the interface that embodies the interaction between the artist and the image making process,
- the interface required for the preparation of the artist's basic tools.

The thesis is primarily concerned with the first type of interface, but obviously some work is required to provide an adequate version of the second.

An artist working with WET & STICKY has to control complex basic materials. These are far more complex than those found in traditional computer based painting systems, but intrinsically no more complex than those the artist will encounter in the studio. It is the manner of presentation that is different. The interface must encourage the user to draw on their traditional experiences and skills, rather than require them to adapt and, or learn new skills.

One has to provide a way for the artist to choose, alter, test and appraise his substrate, and to mix, and modify his paints. This has to be available for use in a natural, unobtrusive manner. To achieve this it seems sensible to refer to how an artist would traditionally accomplish these tasks.

6.2 Mixing paints

To mix traditional oil or acrylic paint the artist will combine quantities of coloured paint from the tube with various mediums. In the studio the artist assesses the paint by its appearance and its feel. The "further work" chapter, discusses how one might provide the artist with equivalent haptic feedback. Assuming we cannot at present incorporate adequate haptic feedback we should concentrate on the process of mixing and its visual appearance.

Typically the artist will arrange his palette with the paints and mediums around the edge leaving the centre as a mixing area. The exact arrangement is highly personal, with some artists always using the same layout so that they always know where to find a particular colour or medium, while others may not have such an organized approach.

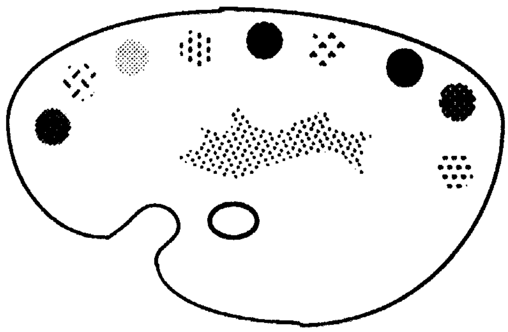

Figure 15
The artist typically arranges his paints and mediums around the edge of the palette, leaving the centre as a mixing area

To generalize the complex mixing activities one should take the simple case where the artist has only two elements to mix, for example paint and medium. Figure 15 shows how an artist might mix these two elements in such a way as to generate a pool of mixed paint that contains a sliding scale from pure paint to pure medium. This is basically a concentration ramp, which can be represented as a slider.

It can be argued that this is a legitimate use of a slider, since it closely mimics the natural control it is representing. More often the association between virtual device and what it is controlling is at best tenuous# . But even in the case of mixing paints what actually happens, when many paints and many media are involved, is too complex to be modelled by any single input virtual device. There are simply too many dimensions for uncluttered presentation via sliders.

An example would be the use of two dimensional arrow icon as a button to traverse the multidimensional space of a hypertext system

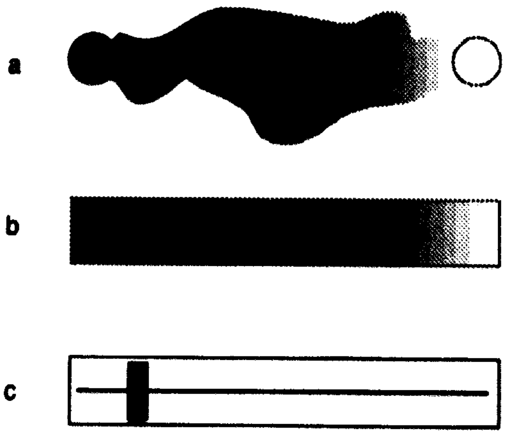

Figure 16
a - illustrates how by the process of mixing paint, black, and medium, white, a pool of different grays forms, b - shows how this may be generalized into a ramp from black to white, and, c - shows that this can be represented as a slider.

When one is only mixing colour, a purely visual attribute, a direct copy of real mixing may be used. An approach used by many computer based paint systems, for example the Quantel™ PaintBox, is to provide a sketchpad, an area where the artist may mix up different colours by dabbing, and blending a small set of initial pure colours into a much wider set of tones and hues. However with non visible attributes like drying rate, and mixability, it is hard to see how to employ the sketchpad method. Although it would be possible to allow the artist to paint onto the sketchpad and alter the values of any attributes to see how this would effect the paint, the problem is that it is impossible to view the results immediately. Unlike colour, altering a paint's viscosity alters its behaviour over time, and the length of this delay is dependent on the paint's other attributes and those of the substrate and other adjacent paints. For this reason the model employs two types of mixing, sketch pad for the mixing of colour, and sliders for the other attributes. Each method is appropriate to its task.

6.3 Selecting and modifying the substrate

In the studio the selection of a substrate is done by colour, feel, its weight, stiffness and texture, and its absorbency. The intelligent canvas supports no notion of weight or rigidity, but does have colour, texture and absorbency. If haptic feedback were available then texture could be selected by the touch and visual appearance of the canvas.

In the studio, to select a substrate by its absorbency, the artist can either choose from a selection whose absorbency he already knows or select a substrate and alter its absorbency#. In either case the artist is effectively choosing from within a range. He may only resort to modifying a substrate's absorbency if there exists no substrate which combines all the other qualities he desires. It is usually easier to alter the absorbency of a substrate than to change its texture. With synthetic canvas there are no constraints on the possible combinations of texture, colour and absorbency. The artist is not reliant on the paper manufacturer producing the desired paper, the retailer stocking it, and there being any left on the shelf.

A simple implementation of the model, devoid of haptic feed-back and a bump mapped display, Blinn (1978), would not include texture as an attribute, because there would be no way of showing it to the artist. All that could be chosen is colour and absorbency. Colour could be chosen from a colour table or as with paints, by mixing. Absorbency could be selected via a slider.

Now the artist has his chosen substrate in his studio. He has now to decide on the orientation of the substrate. Should he place it on the floor, or at an easel, or at 15 degrees from the horizontal (for water colour painting). Sometimes this choice is driven by the desire for specific effects which are the result of a given orientation and other times convenience governs choice.

The user of WET & STICKY also has to make these decisions but it is harder for him to make them in such a direct manner as placing the canvas on an easel. There is a problem with the user's perception of the substrate orientation being linked to that of the screen. It is usually just off the vertical, and it makes sense to think of the substrate being the same. The problem is compounded by the ability of the intelligent canvas to have localized differences in the direction and strength of gravity. Most of the artists I have discussed this with, find this an extremely odd but exciting concept.

One way of setting the orientation of the canvas would be a 3D representation of the plane of the canvas. The artist could rotate the plane around into the desired position. Alternatively a simple slider and compass could control gravity strength and direction. Once the initial orientation of the whole canvas has been selected then the artist is free to load his brush with gravity strength or direction and paint on localized differences.

This can be done by applying a coat of size to the substrate to decrease its absorbency, or in some cases it is possible to increase a substrate's absorbency by washing away its protective coating, or sanding the surface.

6.4 Interface for brushing on attributes

As well as controls for mixing paints and selecting substrates the user needs controls which allow him to load the brush with attributes. Once loaded the term I have used to describe how the artist applies these attributes, is to say that the artist can *paint on* the attributes. This term may be interpreted in two ways,

- that the loaded attributes are deposited on the canvas as independent entites. These isolated attributes then mix with the existing paints on the surface,
- the attributes on the brush are used to reset those of the paint on the substrate or of the canvas.

The former is just a less accurate version of the second. With the first, the user is guessing at the resultant mixed values of the surface paint's attributes, whereas the second allows the user to set any attributes exactly. For this reason it is the second interpretation that is used in the model. To facilitate this the paint selection and brush loading interface requires a method of switching between applying paint and re-setting attributes. In the application mode the user may only apply complete paint particles, with all their attributes intact. When the user is in reset mode then any individual or combination of substrate paint and, canvas cell attributes may be directly altered. This also includes the ability to alter the volume of paint particles held by the cell.

Although the model would support the concurrent application of new paint and the re-setting of existing paint and canvas attributes, there is no need for this added complexity. Again the same results can be obtained by simply resetting existing values. Figure 17 shows a paint selection and brush loading control panel.

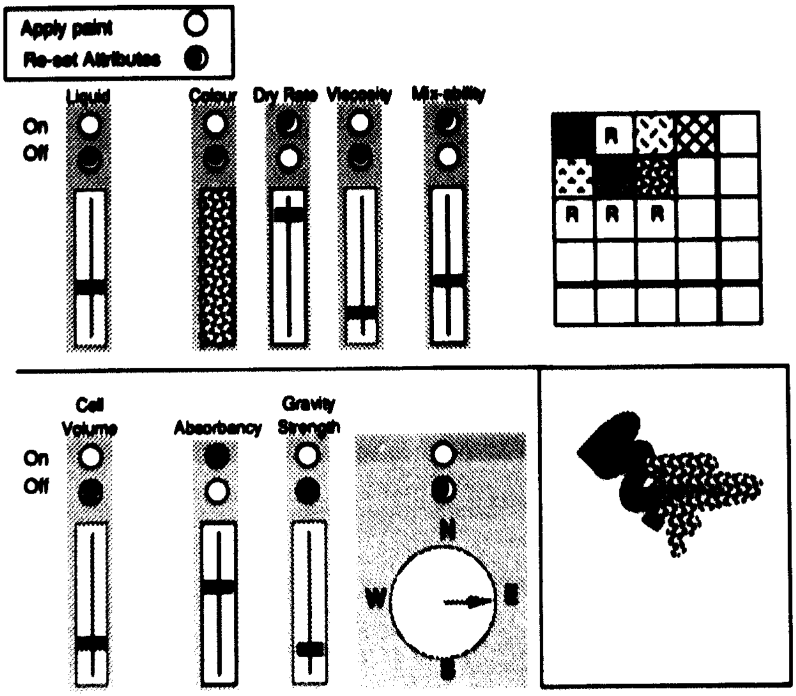

Figure 17
Brush loading control diagram.

The colour mixing and selection controls are absent from the diagram. The area to the lower right is the sketchpad. Just as existing systems use this area for mixing the colour of a paint, so it can also serve as an area where the effects of altering attributes may be tested. The gridded area contains paints that the artist has previously mixed and wishes to retain. An R indicates that the paint is a reset paint, rather than apply paint. On selection of one of these paints the controls would adjust the display to show the paint's qualities. The controls below the horizontal line are all controls associated with the substrate, while those above are associated with paint.

If the *apply paint* toggle were on, then all the substrate controls would be greyed out, indicating they were inoperative. If, as in the diagram, the toggle were set to reset attributes, then all controls would be potentially operable. The user could set which attributes he wished to load onto his brush. Using the brush in this mode he could carry-out very specific operations upon the image. In the illustration the brush is loaded with ; dry rate, mixability, and absorbency, which causes

- drying up patches of paint,
- altering the ability of the paint to mix with others,
- a reduction in the liquid content of any paint applied to that area of the substrate

6.5 Input

In the present implementation the painting engine performs one evolution of a cell and then looks for any user events. If it does not detect any then it selects another cell and performs another evolution. If there is user input in any form then the painting engine is temporarily halted while the system deals with the user's events. These events can take the form of selecting paint from a pre-mixed pot, mixing new paint, changing brushes, or applying paint or resetting paint or canvas attributes.

In the case of applying or resetting paint or canvas attributes the system updates the affected cells in either one or two passes. If paint is applied to the canvas then two passes are used. The first pass changes the colour attribute to any cells affected by the stroke to that of the applied paint and keeps a list of all such cells. When the stroke is finished the second pass uses the mixing section of the painting engine to mix the applied paint with that already on the surface and to update the canvas cells and their contents. In the case where the user is resetting attributes the system updates the paint or canvas attributes in a single pass. A single pass is possible because there is no need to incur the delay of having to invoke the painting engine. In both cases, once the user input is completed the painting engine returns to its normal task of evolving randomly chosen canvas cells.

6.6 Impersonating traditional systems

If a user only applies paint which has zero liquid content then the model behaves in exactly the same way as traditional systems. The painting engine visits the cells and upon discovering that they are dry, makes no attempt to move the paint by diffusion or dripping. In this way the WET & STICKY model can re–create any of the effects possible with traditional systems based on the Shoup model. It is an important point to note that the Shoup model is a sub-set of the WET & STICKY model.

It is possible for the WET & STICKY model to make use of all of the tools and effects that have been developed for the Shoup model. With the addition of an interface similar to the one in the previous section any tool could be used in its normal manner, but rather than only using colour as input it could be used to apply paint, reset paint or reset specific attributes. Arithmetic effects traditionally applied to colour values could be expanded so that they could apply to any one or more of the paint's and substrate's attributes.

6.7 Conclusion

This chapter has explained the requirements of an interface to a version of WET & STICKY. It has shown that a more complex interface is required than for normal systems based upon the Shoup model. Although this thesis presents no evaluation of the interface proposals, this would form an important part of future work. This would entail investigating the best ways to map users existing skills and experience onto the new environment, and to build an interface which allowed access to the power of the model without alienating the user.

Conclusions

Chapter 7

7.1 Summary

After the process of implementing a version of WET & STICKY, it is necessary to reappraise the abstract model. The implementation casts doubt on some aspects of the simplicity of the initial model, particularly the formation of drips and the requirement to model surface tension. This section shall expand the initial model, re-building those areas that were flawed or not sufficiently understood at the outset.

7.2 Reappraisal

WET & STICKY is more computationally expensive than the relatively simple Shoup model. As the implementation progressed it became obvious that additional parameters over and above the initial model were required to achieve realistic results.

The major additions were surface tension, diffusion and the stochastic methods employed by the painting engine. Although with hindsight it is possible to say that initial model was lacking in that no attempt was made to include surface tension and diffusion, it was the process of experimentation with simplistic models that highlighted the level of complexity that was required. The initial attempts at surface tension algorithms were complex baroque structures, which through experimentation, redesign and distillation have resulted in the elegant and simple solution described in Chapter 5. The gravity strength algorithm is also a simple and effective method of modeling of a real world effect. However all these additions increased the computational work load with the result that the present implementation evolves at too slow a pace to make it a viable interactive tool.

The whole process of the implementation described in this thesis can be likened to the construction of a tower with children's building blocks. The final structure is the result of reapraisal and re-construction of many failed attempts. As one builds tier upon tier, the stability of the whole increasingly depends on the integrity of the foundations. The foundations supplied by the initial model were not strong enough to support the planned structure. The implementation described here is a sound structure, but it is not advocated as the only way to build a structure upon the strong foundations of WET & STICKY.

The images in this thesis are included to show that the model does allow the paint to evolve, that it is possible to set independent paint and substrate attributes, that drips form, and that mixing occurs. They do not show evidence of interaction with the system. The images are taken from an implementation carried out by two researchers in this department, Dr David England and Dr Kevin Waite. Their implementation was derived solely from the specifications provided in chapters four and five of this thesis.

At present the main problem of WET & STICKY is the speed of interaction. This is entirely due to the complexity of the model, the depth of information required for each canvas cell and collection of paint particles, and the large amount of data modification that is required for each cycle of the painting engine. Before WET & STICKY can be realistically considered as an interactive painting medium the speed of the system will have to be improved. There are several way that the speed of the system could be increased.

An attractive improvement is to have a wet list, that is a list of cells which contain wet paint and only have the painting engine visit the cells on the list. However the problems and overheads in maintaining such a list and the complexity it introduces into the implementation outweigh any speed advantages. It was found that the speed advantage derived from a wet list was lost after about 25% of the canvas was wet. After this point there was no speed advantage and as the proportion of wet cells increased the maintenance costs made a wet list a positive disadvantage.

The present version makes use of the SunView window manager and its graphics procedures. If the system were designed to support its own windows and could write directly to the screen then it could potentially run faster. With the benefit of the experience gained writing a development system, a more ordered and efficient second generation system could be produced. Although it is unlikely to happen, further speed enhancements could be made by implementing the system in machine code.

Another approach is to adapt WET & STICKY to run on a parallel machine. At present Dr David England is testing different processor configuration to determine which provides the best compromise between speed enhancements and data-passing bottle-necks. So far speed improvements in the order of 100% have been achieved when using a 8 processor array#. The reason this is such a small speed increase is partially that each processor is running at 7 mips, as opposed to 12mips for the Sparc Sun, but also a drawback of the current graphics implementation. The graphics package used (running under the CS Tools system) only transfers one pixel at a time between the computing and graphics

Meiko Computing Surface

processors. This does not make full use of the inter–processor communications. Thus the screen is not up–dated at the optimum rate.

More work will also have to be done to test the effectiveness of the present interface for the selection and modification of paint and substrate attributes. Despite these problems I believe that this thesis successfully demonstrates that WET & STICKY provides a novel model of electronic painting, and one which allows the creation of rich and complex marks. Furthermore, by basing the model on the real world media and environmental effects, and by abstracting these into individually accessible and definable entities, it is possible to use WET & STICKY to create effects which previously were impossible.

If we return to the self imposed criteria for success defined in the section, "Stage-set verses reality", then I believe that the system is successful. Paints appear to evolve according to the user's expectations of how, in reality they should evolve. The resultant images possess a richness and complexity of marks that would be very difficult to produce with existing systems. This richness is not necessarily the result of explicit actions by the user, but of the continual actions of the painting engine. However as yet no real evaluation has been carried out with fine art users to validate these claims but this will be an integral part of any future work.

In the introduction to chapter one a promise was made that this thesis would find a cure for the ills of existing computer-based paint systems. I believe that the WET & STICKY model does fulfil this promise, and that the photographs included in this thesis provide proof . However before it can be usefully applied the cost of the treatment, in terms of the speed disadvantages, will have to be solved.

Colour Plates

This section includes a selection of colour photographs of the effects possible with WET & STICKY. They are intended to show the affect of varying different attributes. All photographs are of 8 – bit deep images.

The first image is a control which shows the evolution of nine square blobs of default paint. The default values are,

paint liquid content	80
paint drying rate	80
paint mixability	80
canvas gravity strength	10
canvas gravity direction	south
canvas absorbency	10

All values except gravity direction are in the range 0 – 100. Gravity direction can be either north, south, east or west.

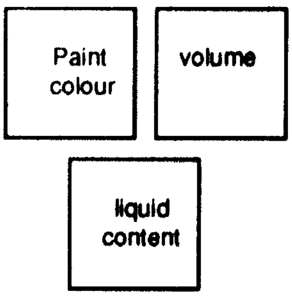

If only one window is shown it is an image of the colour of the paint. If two are shown then the left image is of the true colour of the paint and the right is a volume map, that is, an image of the volume of the paint held by the canvas cells. If there are three windows then the left image is the paint colour, the right is a volume map, and the bottom image is a liquid content map. These alternative views are included as they give a better impression of the formation of drips and the mixing of paints.

In all the images illustrating the effect of altering paint or canvas attributes, the colour of the paint is derived by using the value of that attribute as an index into the colour table.

The no callibration of the attribute values has be carried out on the implementation from which these photographs are taken. Some of the effects shown are stem from this lack of callibration. For example, although the range of a value may extend from 1 – 100 the effective range may only be between 30 and 70.

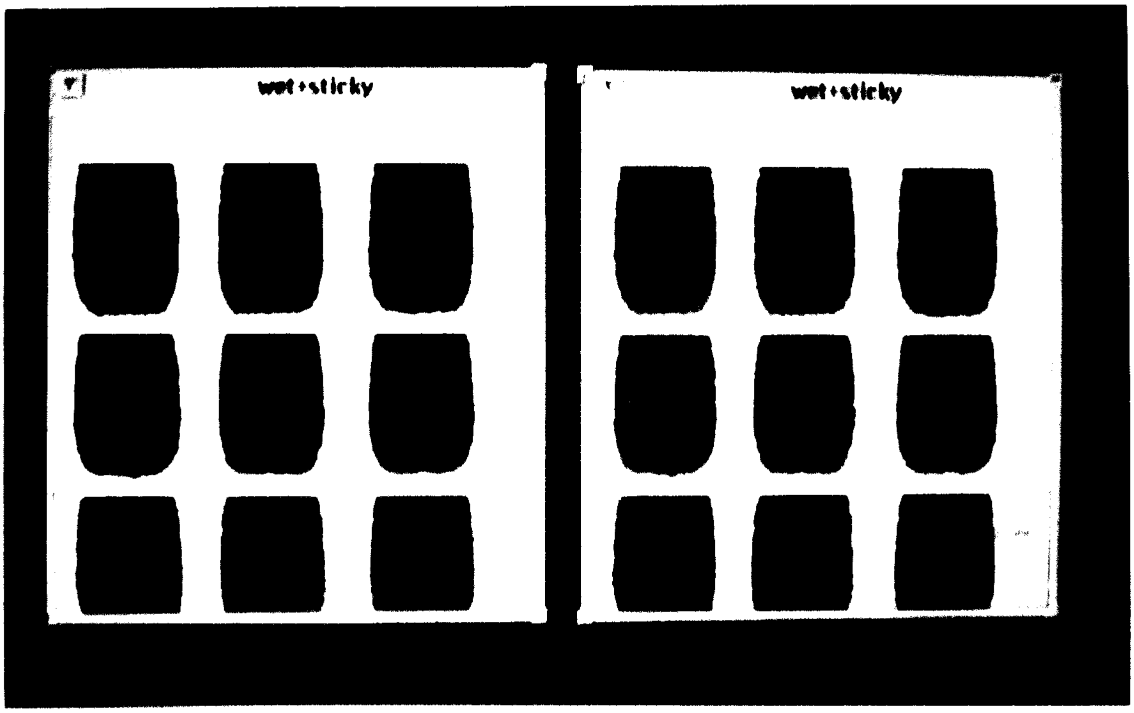

Plate 1 Control image

The evolution of nine square blobs of default paint.

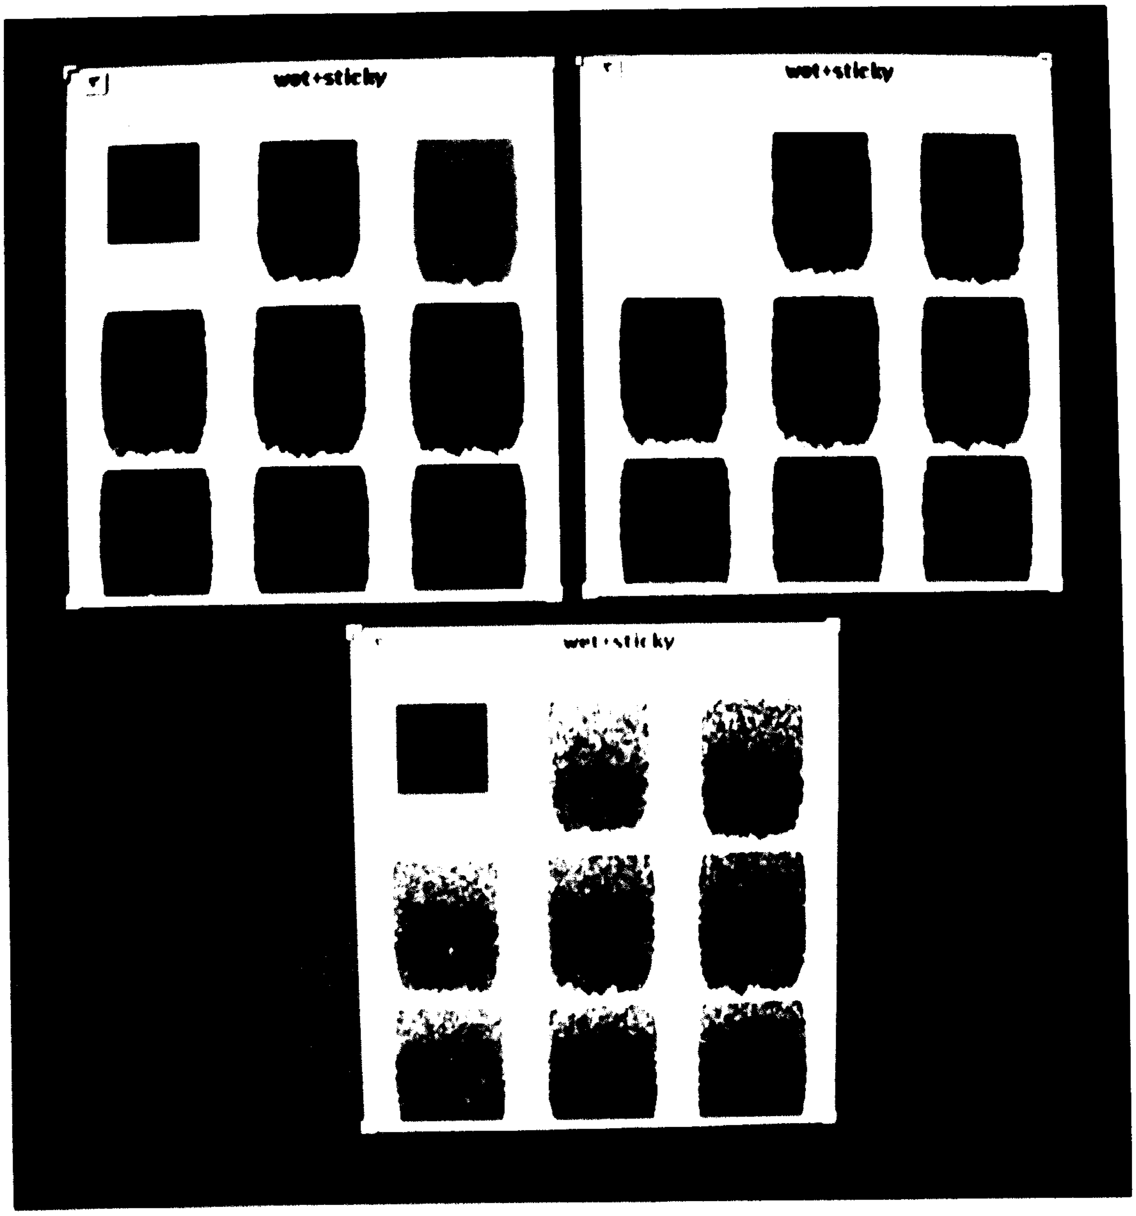

Plate 2 Volume

This image shows the effect of altering the volume of paint held by the canvas cells. The image started as nine square blobs of paint of different uniform volume. Their volumes ranged, in regular steps, from 0, for the top left blob, to 100 for the bottom right. All the other paint and canvas attributes are set to their default values. The colour ranges from white representing 0, through blue, purple, and brown, to red representing 255.

In the volume map, the top left blob is invisible because it contains no volume. The volume map also clearly shows the increase in volume, and the progressive thickness of the drips that are formed, across the other blobs. The liquid content map shows that all but the top left blobs are nearly [illegible] The top left blob has not evolved and has remained with its default liquid content of 80.

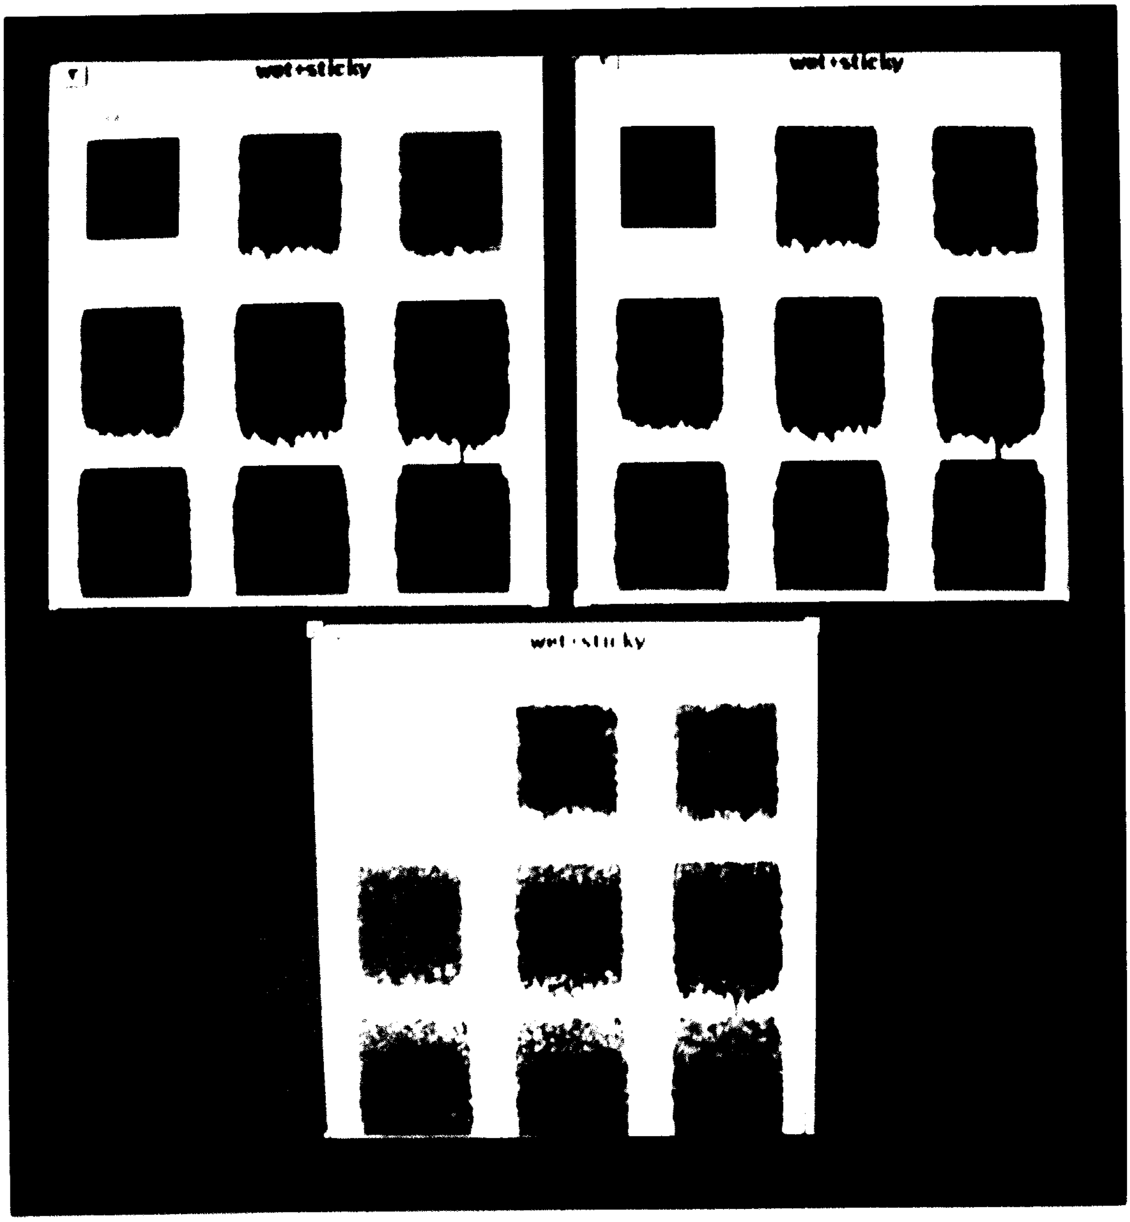

Plate 3 Liquid Content

This image shows the effect of altering the liquid content of paint. Again the range starts from 0 and increases to 100. All other attributes are set to default values. In this case it is the liquid content map that shows the top left blob as having 0. Altering the paint's liquid content affects it rate of flow. This is evident from the volume map. Here the top left blob of dry paint has not moved at all, whereas the bottom right, which had paint of maximum liquid content, has moved the furthest. The bottom row of blobs have formed pools of paint where they have reached the edge of the canvas.

The liquid content map shows that thin trailing edge of all the drips is dryer than the thicker leading body.

Plate 4 Drying Rate

Here we see the effect of altering paint's drying rate. In this case the drying rate starts at 0 and increases to 100. All other attributes are set to default values.

Each time the painting engine visits a cell which contains wet paint it uses the paint's drying rate as a measure of the probability that it will reduce the paint's liquid content. The top left blob is made of paint which never dries whereas the bottom right always dries.

The only limit to the movement of the top left blob is its original volume. This is because the absorbency of the canvas is not zero and therefore it retains a number of paint particles. This can be seen in the volume map where the top left blob is very thin. It can also be seen that he middle left blob is in the process of loosing its thick leading as it flows into the blob below.

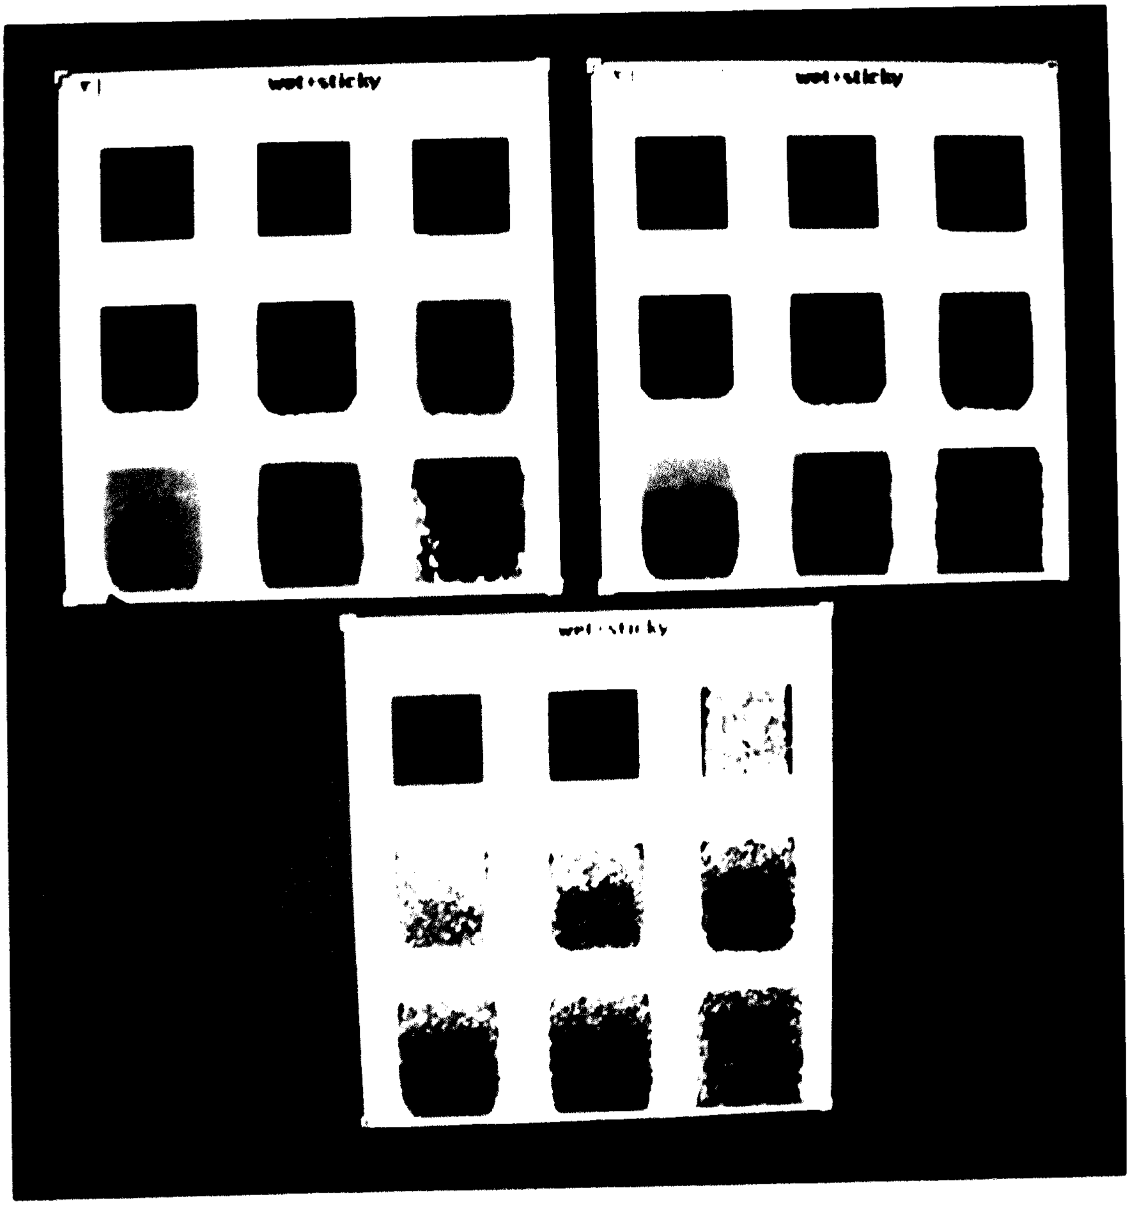

Plate 5 Absorbency

This image shows the effect of altering the absorbency of the canvas. The area under the start locations of the blobs of paint ranges from 100 in the top left to 1 in the bottom right. All other attributes are set to default values. The effect of a cell's absorbency is to retain paint particles within that cell. In the volume map we see that there is little or no movement in the first three blobs, with progressively more in the other blobs. The colour of the first two blobs in the liquid content map are the result of an artefact of the colour table

BEST COPY AVAILABLE

PALE PRINT ON PHOTOS

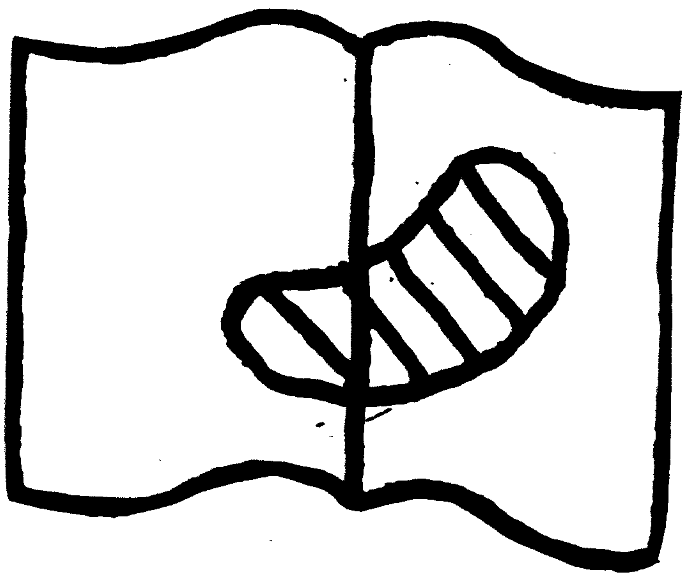

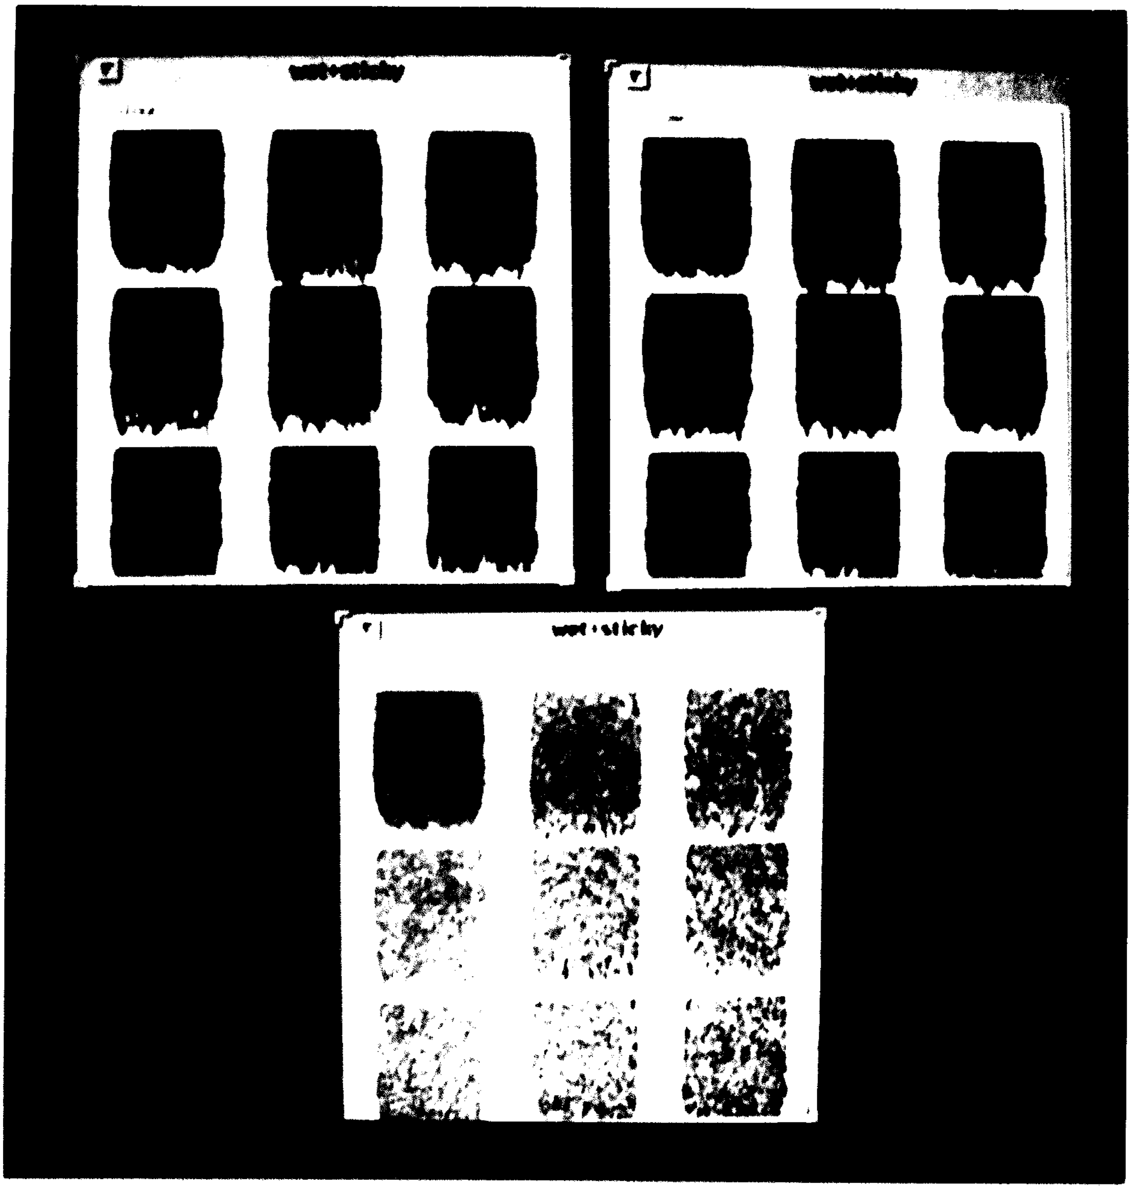

Plate 6 Gravity Strength

Changing the gravity strength has an effect on the ability of paint to pass from one cell to another. Higher gravity strengths means an increase in the gradient of the rear slope of a drip. This in turn results in an increase in the volume of paint held by cells in the body of the drip.

The this example gravity strength ranges from 0 in the top left blob to 100 in the bottom right. In the volume map the top left blob can be seen to have moved down a short distance. This is due to the fact that although there is no gravity acting upon the paint, diffusion is having effect. Diffusion only acts upon the bottom , left and right sides of a cell and therefore the paint has spread down and side ways.

The results for the other blobs require more explanation. Firstly the changes made to gravity strength only applied to the areas under the start point of the blobs. This has meant that once paint moved out of these areas it has behaved as if it were default paint upon a default canvas. This explains the start of the formation of drips of standard appearance. Secondly, the increase in gravity strength has resulted a greater proportion of paint being passed by gravity than by diffusion. This is evident in the volume window where one can see that as gravity strength increased the size of the red thick patches of paint decreased. This shows that columns of paint have formed. An additional problem is that the thickness of paint has exceeded the maximum that can be displayed by the colour table.

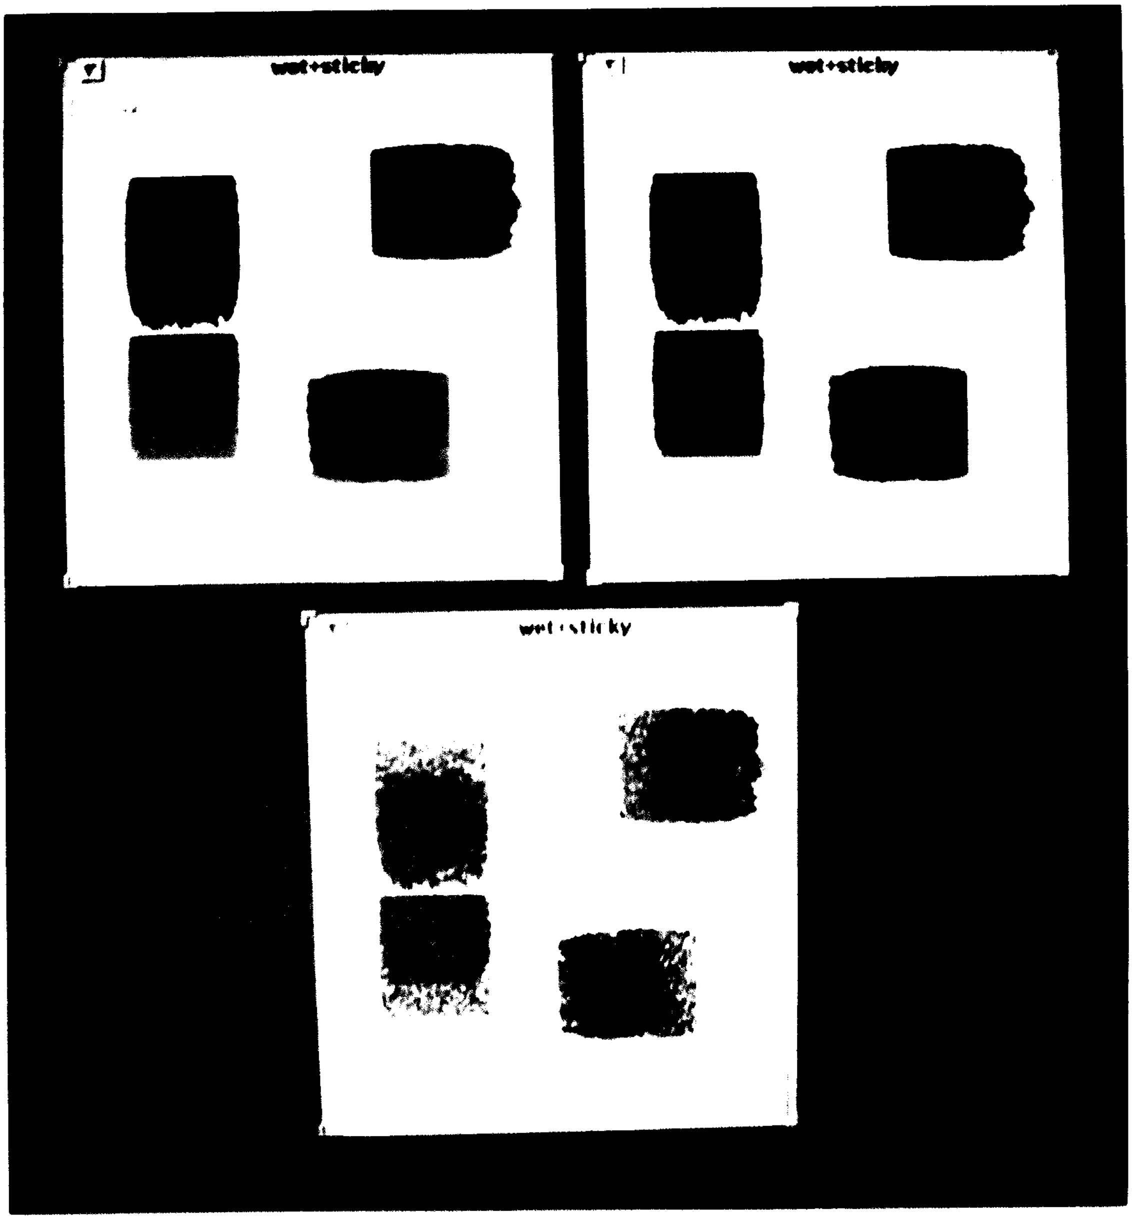

Plate 7 Gravity Direction

In this image all the attributes of both the canvas and paint are set to their default values, except gravity direction. The canvas has been divided into four areas each with a different gravity direction. Upon reaching the edge of its quadrant, the blob travelling north has started to collect in a pool.

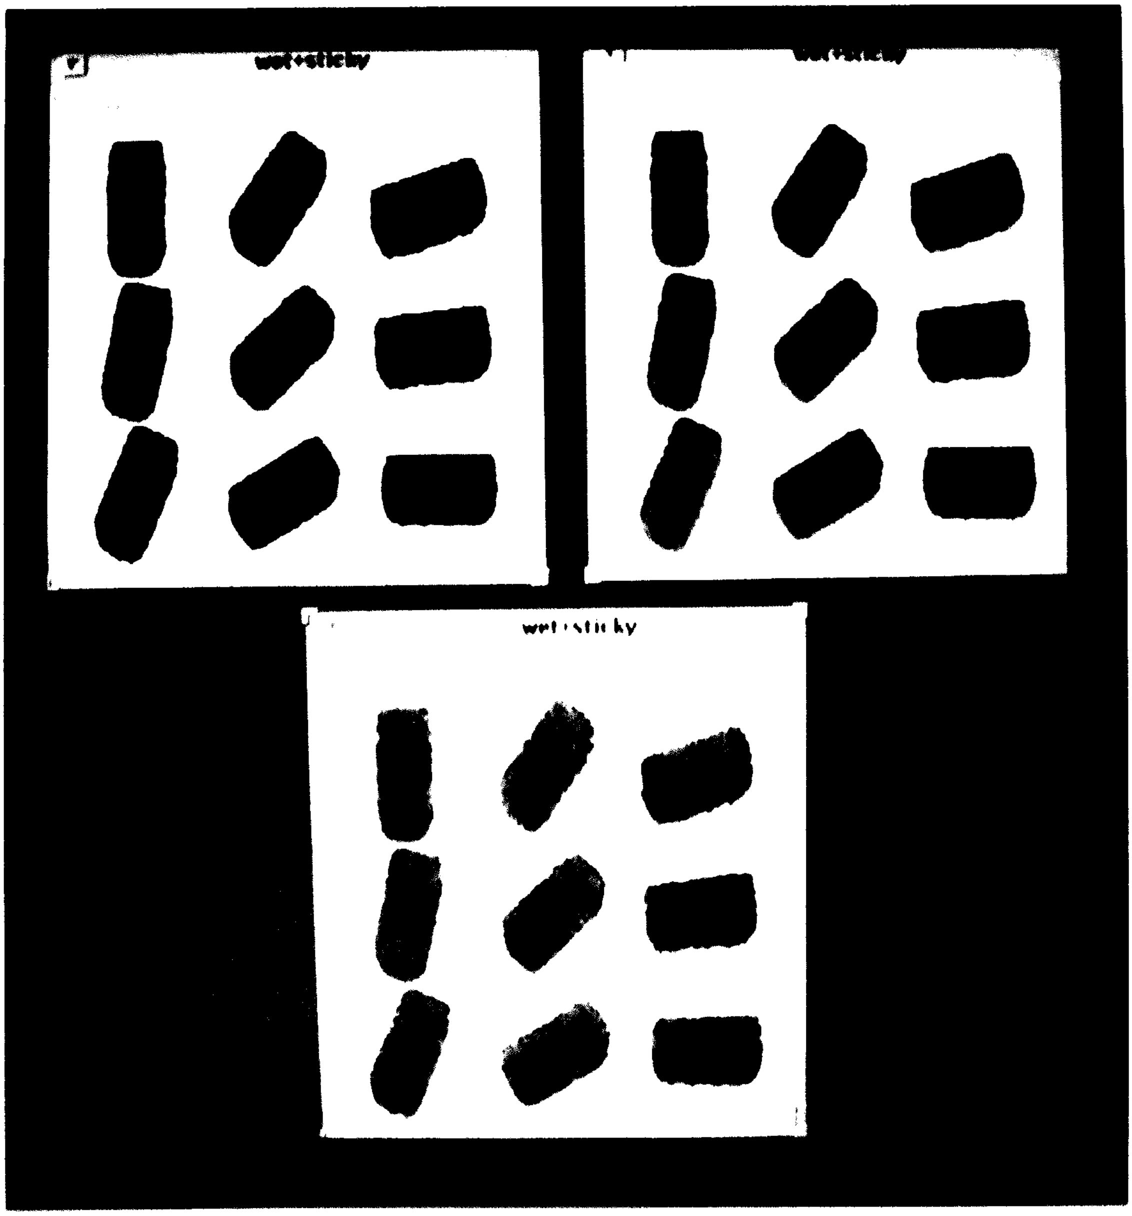

Plate 8 Nine Strokes

This illustration shows the affect of altering the gradient of the edges of a stroke of paint. All the nine blobs started as a rectangle of the same size and have been placed on the canvas in orientations ranging from, vertical to horizontal.

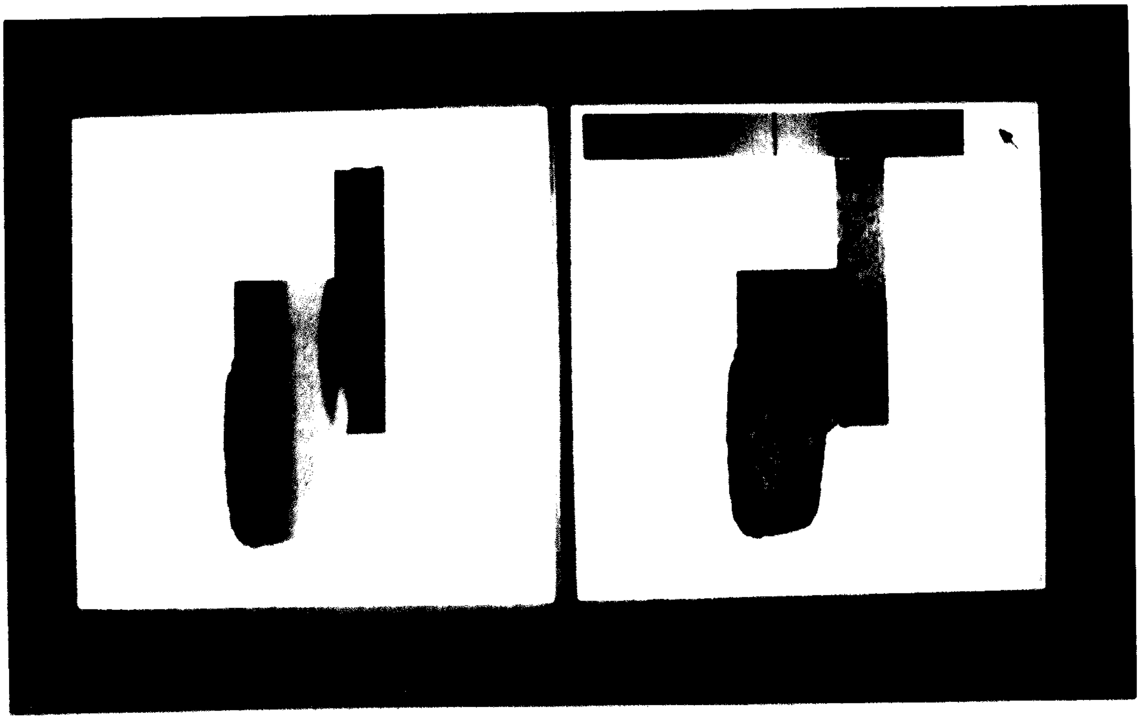

Plate 9 Mixing

This image shows the evolution of three blobs of paint, each of which started as a rectangle of equal size. All three paints are identical except for their colour. The red blob is placed upon a strip of canvas which has a gravity direction of north, whereas the rest of the canvas has a gravity direction of south.

Soft mixing has ocured between the blue and green paints. The swirl mixing that has ocured between the red and the green is a result of diffusion passing paint across the gravity direction line, so that red paint that was travelling north has mixed with green paint travelling south, and vise versa. It is also interesting to note the thick shelf of paint that has collected at the top right of the green blue drip.

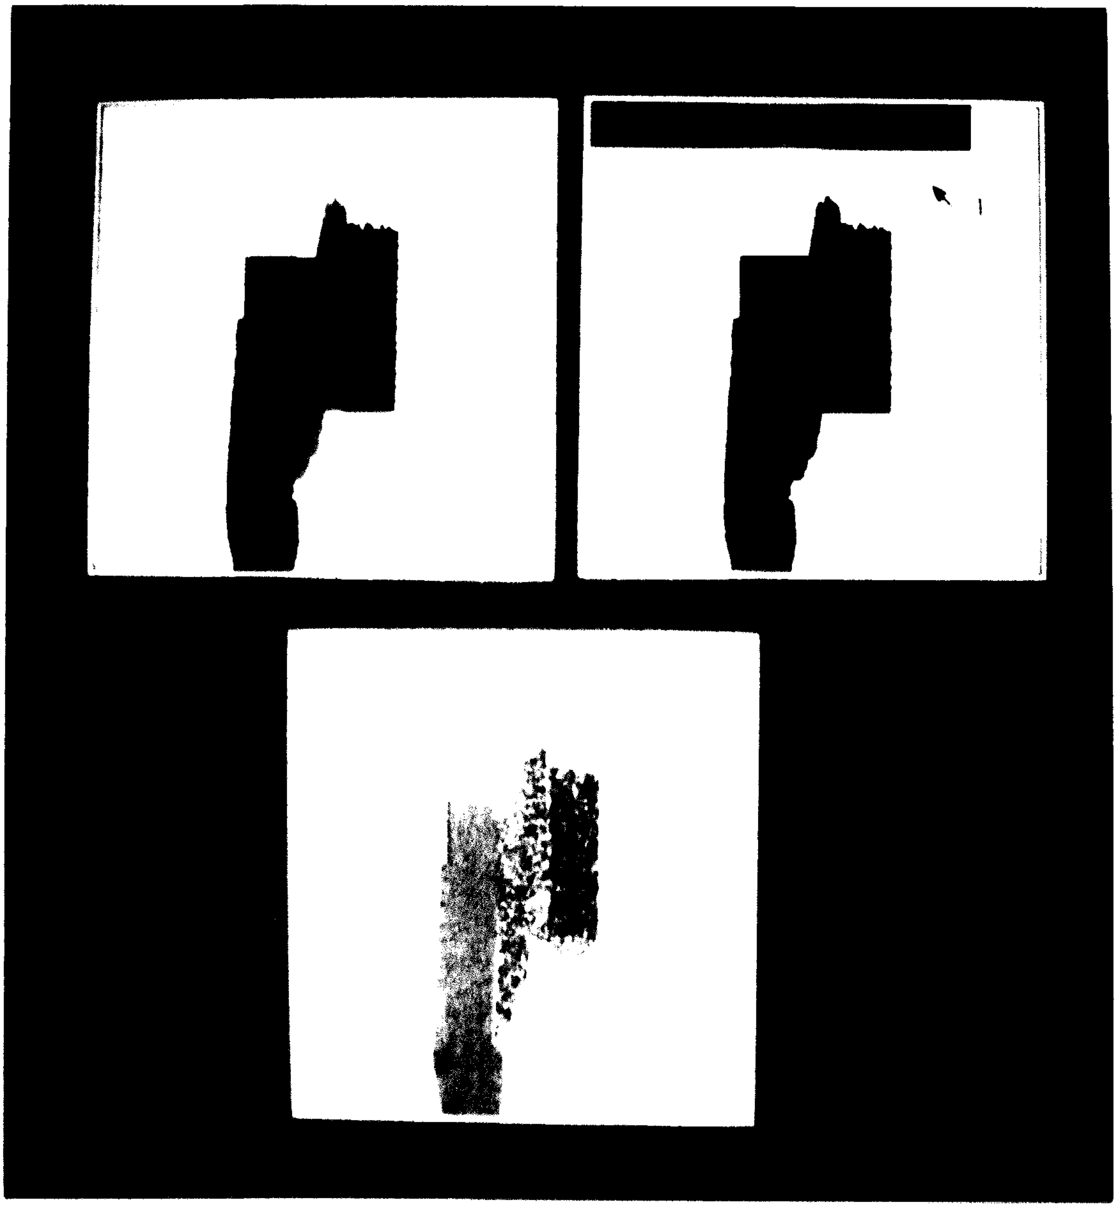

Plate 10a Mixing of three types of paint

The next two images illustrate the evolution of three different types of paint over a period of time.

The values of the paints are,	blue	green	red
liquid contents	100	80	50
drying rate	10	50	80
mixability	80	80	10
volume	50	30	50

All other attributes except gravity direction are set at their default values. The left half of the canvas has a gravity direction of south and the right half of north. The images show that the slow drying, very liquid, blue paint keeps flowing after the other two paints have dried. It also shows that the more viscous red paint does not travel far as the others and that there is very little diffusion from the red paint onto the bare canvas.

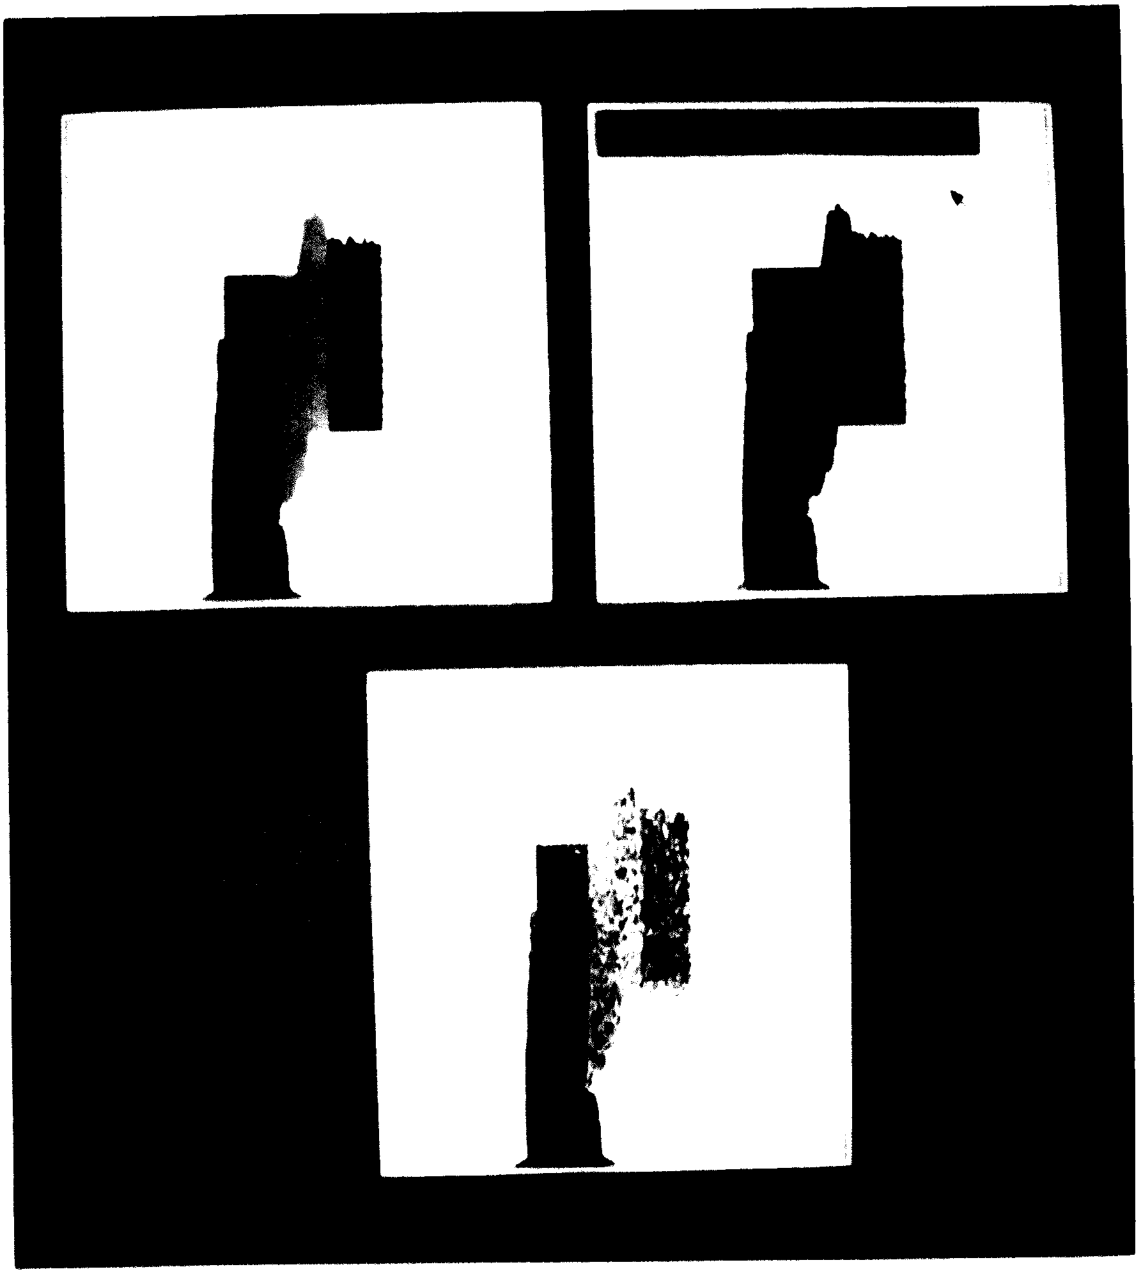

Plate 10b Mixing of three types of paint

It can be seen in this image that the blue paint is still not as dry as the green or the red, and that it has collected in a very well formed pool at the bottom of the canvas.

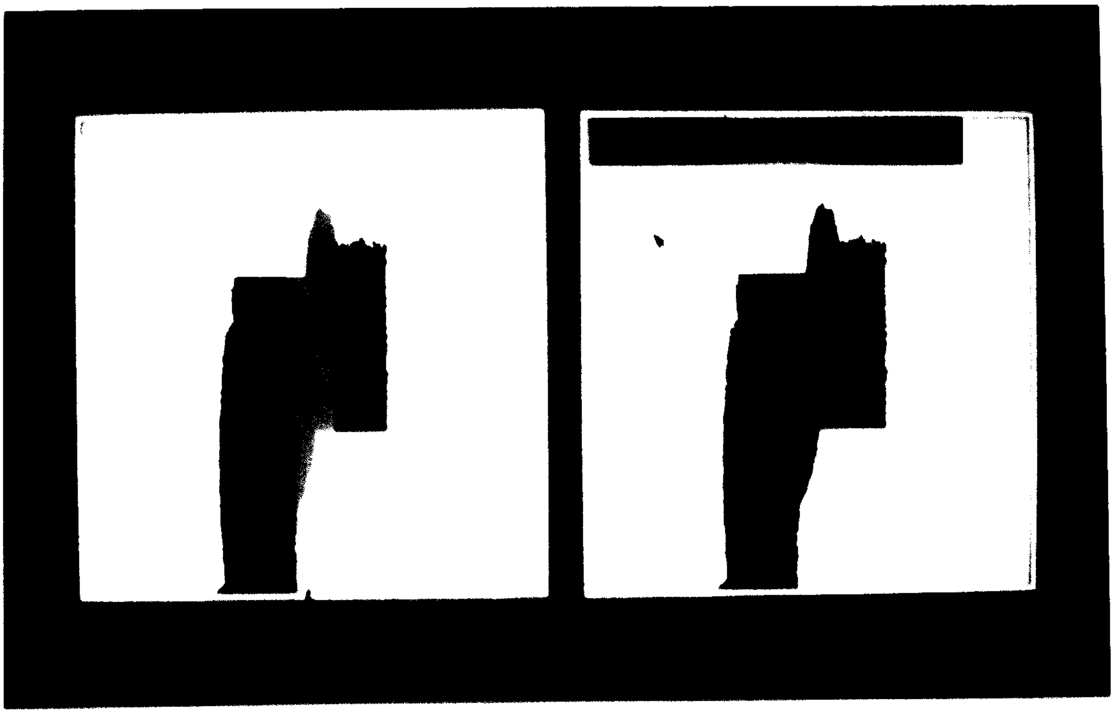

Plate 10c Mixing of three types of paint

This image is included to illustrate the relative unpredictability of the WET & STICKY model. The paints and canvas attributes are exactly the same as in the previous example but the drips and mixing that have occured are not the same. This image was taken after the canvas had been evolving for the same period of time as the image in plate 10 b

Further work

Chapter 8

8.1. Introduction

This section describes advancements and additions that could be made to the model. These can be devided into two basic areas, user input, and the complexity of the model. I shall first discuss advancements that are required for the accurate simulation of brush strokes and then go on to describe systems which possess greater levels of complexity but which can nether the less be derived from the basic WET & STICKY model.

8.2. Brushes

The most accurate and complex brush simulation thus developed is Strassman's (Appendix B). Although it is major improvement over traditional "potato print" brushes, it is still lacking in several ways. It was designed to work with paint systems based on the Shoup model, and therefore only deposits colour. It can be modified to deposit paint rather than colour but it has no notion of brush dynamics and rules governing how these affect any existing paint on the substrate.

Strassman only models the effects of change in pressure applied to the bristles of the brush, the path of the stroke, and the load capacity and distribution of paint amongst the bristles.. Ideally to provide a more complete model of brush dynamics one should also take in to account the effects of the tilt, rotation, and acceleration of the brush, and the flexibility of the bristles. These would combine with the properties of the paint on the brush *and* those of the surface paint *and* the properties of the substrate itself. For example; the viscosity would influence how the paint flowed off the brush, while the absorbency of the canvas would have an effect on the length of a stroke.

The bristle flexibility would affect the behaviour of the brush, particularly when there was an abrupt change in the direction, or acceleration of the stroke. Under a sudden change of direction, stiff bristles loaded with liquid paint would splatter paint in the original direction of the stroke. This effect would be less likely with a softer, more flexible brush or with more viscous paint.

A wet brush moving over wet paint mixes with, and picks up, the surface paint. Some brush strokes plough through the surface paint, others glide over. In the studio the artist has the option to use a brush that is contaminated with more than one type or colour of paint. He may use a dry brush simply to move the paint about, or to cause it to mix[#] .

Again it is the accidental and unpredictable serendipitous effects that should be incorporated into the simulation, and again these should be based on empirical observations. The use of a more responsive and informative input device is also a prime requirement of any complex brush model. If the image of the painting is displayed as a bump-map as described earlier, then the user receives visual information about the qualities of the paint on the surface. If a normal 2D image is presented then the user receives much less information (although the provision of attribute maps would allow the user to see the non-visable attributes).

Ideally as well as being able to see the effects of the texture of the canvas and paint, the system should allow the user to *feel* the texture, and *feel* the viscosity of the paints on the surface. In the studio the artist uses the haptic channel i.e. touch feedback, as a way of gauging a paint's relative viscosity or liquidity. The resistance felt by moving a brush across the canvas gives an indication of both the qualities of the surface, the paint on the brush, and of the paint already on the surface.

Haptic feedback also plays an important part when the artist mixes a new paint. In this case the artist would add a little more water, or medium to make the paint more runny, or to thicken it, and use the feel of the paint as a primary way of gauging its consistency. So some form of haptic feedback should be incorporated into the model. All the information required is already present, all that is needed is a way of interpreting and relaying it to the user.

[#] This is essence the effect modeled by the smudge and smear tolls provided with most paint systems. However these tools, working a simplistic model, can only really give the effect of causing the mixing of wet, water colour paints, and not the range of effects that would be possible with the Wet and Sticky model.

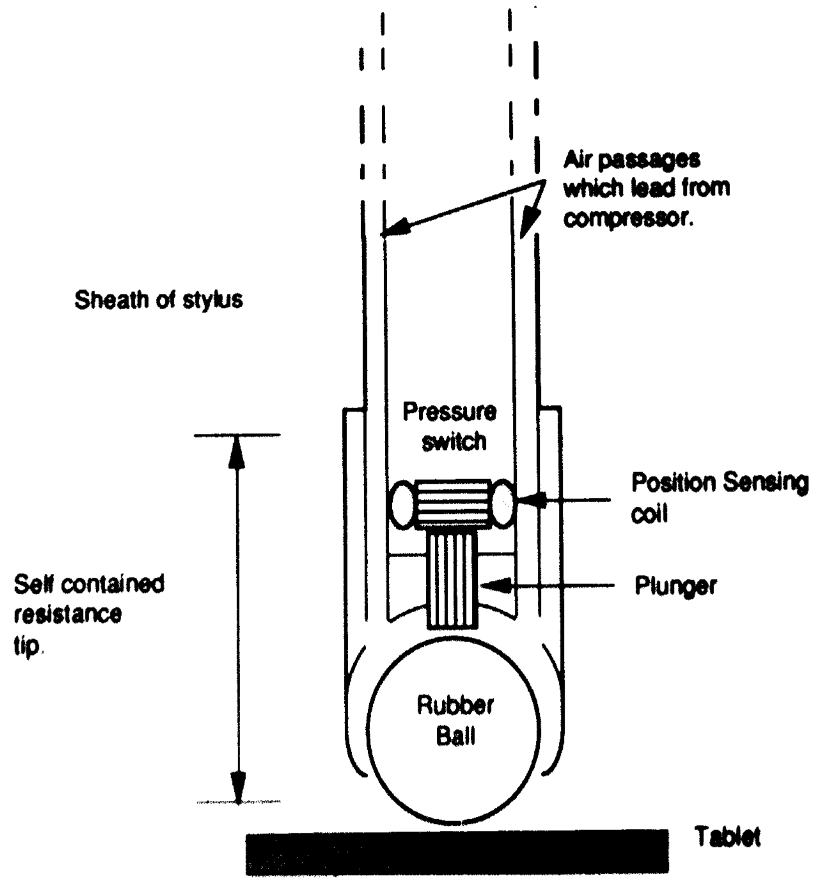

Figure 18
Illustrating a possible design for a reactive stylus.

A possible solution to this problem would be to have a reactive stylus. As yet no such device exists. A simple version could be made by fitting some form of resistive tip to a normal stylus. Such a tip could provide resistance by using a rubber ball held loosely in a case. The ball's rolling resistance could be influenced by positive and negative air pressure. Air pressure could be supplied by a small compressor of the type used to power air-brushes. Jets of air could lessen the rolling resistance by blowing against the surface of the ball which would make the ball float in a cushion of air. To increase the resistance suction could be applied which would tend to hold the ball firmly within its case. To enable the application of non-uniform resistance three or more pairs of air passages could be used. Each pair would consist of one air passage for positive resistance and one for negative. These would be placed equally around the top surface of the ball. With such a system resistive feedback could be provided while the artist mixed his paints to a desired viscosity and also while he painted over the existing paints.

It seems clear that a similar stochastic approach as that employed by the painting engine coupled with haptic input could provide an effective brush simulation. Work on this area was not considered to be part of this thesis. The author believes it to be an area of similar if not greater complexity to that of modeling paint and canvas, and is an important area requiring a great deal of further work.

8.3. Further Improvements

The model as outlined in the previous chapters constitutes a basic system. The important characteristics, ie. those that differentiate it from the more complex models are,

- the logical distribution of attributes based upon those observed in reality,
- the use of a set of painting engine rules that is likewise derived from reality.
- implies that those attributes associated with paint are those one would expect to be associated with paint.

The artist would interact with the system via a brush, painting on not merely paint but also using the brush to reset/paint on attributes. For example, if a drip was moving into an area where the artist wanted to keep clear, he could halt, or restart its progress by painting on dryness, or greater absorbency, or viscosity, or decreasing the strength of gravity, or altering its direction.

The system would also support the standard paint system tools such as editing, erase, basic drawing primitives, etc. It would also include the search and replace tools such as flood fill, and a more complex version of the basic rasterops. The power of these tools would be greatly enhanced when used with a WET & STICKY system. With standard systems based on the Shoup model the user only has the colour of a pixel as input. The user would select a pixel with a colour value he wished to alter, which we shall call the target value. He would then select a replacement colour value, and apply a tool which seeks out pixels with a target value and resets them to the replacement value. The tool may work like flood fill and work out from a seed point, or may replace all targets within a users defined area, or the user may paint on the area he wishes to be affected.

With WET & STICKY, the artist may use any of the paint's or canvas's attributes as targets or replacements. So rather than simply being able to change all red pixels to blue, he could, for example, make all dry canvas cells become wet again, or, make the paint in cells with an absorbency value of X become highly viscous, or, increase the gravity of all cells containing red paint, etc.

This could be further enhanced by allowing conditional or compound target and replacement values. For example, the target could be defined as all cells that contain red paint and have a gravity direction of north. The replacement could be specified as a completely different paint or may only change one or more of the cell's and/or the paint's attributes.

8.3.2 Configurability

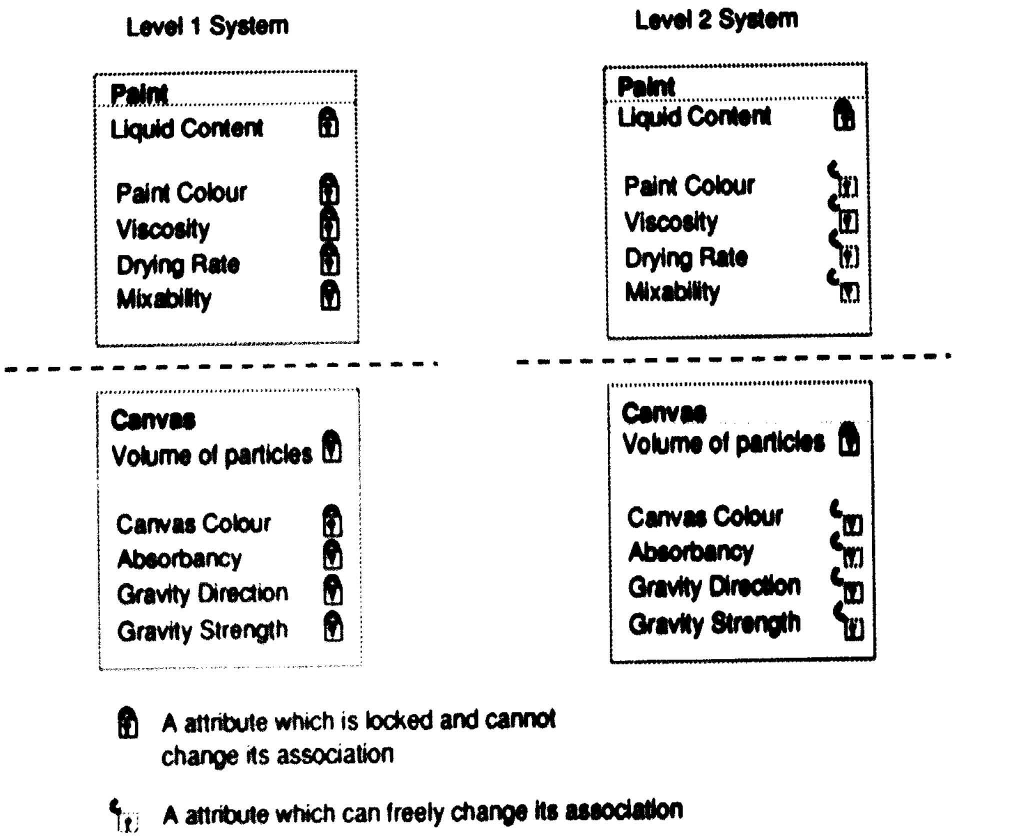

Figure 19
Illustrates the ability to dynamically re-configure the associations of attributes.

The next step up in terms of complexity would be to allow the user to dynamically re-configure the data structure. At any time during the painting process the user could alter the association of the attributes.

In the basic system the constituent parts of paint are those that one would expect to find in nature. In a configurable system, the user could use paint that contained canvas attributes, and canvas that had paint attributes. For example the user could define a gravity strength paint. This is different from using a brush to reset the gravity strength values of the canvas. A gravity strength paint would mean that gravity strength would flow with the paint, collecting in pools and drips in the same way as normal paint.

This re–assignment of attributes could be done across the whole canvas, or on an individual, cell by cell, basis. In the case of the latter, if a cell containing a paint bound attribute transferred paint to one where that attribute was bound to the canvas, then the paint would deposit the attribute in the cell. In this way drips containing an accumulated quantity of an attribute dropped into individual, or areas of, cells creating hot spots of that attribute.

The only attributes which could not switch their allegiance, are, in the case of paint, its liquid content, and in the case of the canvas, its volume of paint cells. The painting engine relies on liquid content to determine the flow of paint, and canvas cells need some record of how much paint they contain. As far as the painting engine is concerned, so long as the paint is liquid and has a volume, it does not matter whether the other attributes are associated with the canvas or paint. As long as it has access to the attributes it can perform its task.

8.3.3 User–Definable Rules

The next level of improvement requires a relatively small conceptual step forward but a great leap in terms of the advancements required from the user interface.

This would allow the user to define his own set of rules for the painting engine. Just as in a data configurable system, the association of attribute may not mirror nature, so in a rule configurable system the rule may not mirror those found in nature. And, following the cellular philosophy of the model, these rules may differ from cell to cell, (Cook 1984). The rules would be painted onto the surface. However these rules could not migrate and could not mix with one another.

Such a system would allow the user to rebuild the painting engine, and construct alternative paint realities, (Smith 1987). Basic and configurable systems would only allow the user to alter the data that fuelled the engine. The engine itself would remain intact and unchangeable. But in a rule building system, the user could use some form of visual language to plug together rules, effects, and associations. This would require a very flexible and powerful interface.

One of the problems with such a system would be that it would lose the advantage, inherent in reality–based systems, i.e., of employing the artist's previous knowledge. In the previous two models the artist has a basic idea of how paints will behave. He can predict their basic behavioural tendencies without needing to have explicit knowledge of the rules employed by the painting engine. Although the system would have the potential for the construction of complex rules it is unlikely that this would happen. I would expect it would be used for the development of fairly simple rules. In this case the system would be behaving as a more generalized simulation tool.

One such new set of rules has already been proposed. Patterson's (1990) new painting engine is envisaged as an aid to the accurate drawing of freehand curves. He proposes a much simpler version of the intelligent canvas and paints with a reduced set of attributes. With the passage of time his engine would increase the transparency of lines applied to the surface. If the user were to repeatedly sketch over the path of a desired curve, then only the central or common elements of the sketch sub–curves would remain dark. This darkened area would then be taken as the desired path through which the engine would compute and display a cubic B-spline.

8.4. Additions to the basic model

8.4.1 Enhancements to the Interface: Alternative Views and Named Areas

Since many of the paint and substrate attributes are not normally visible, a future interface should supply alternative views onto the image.

In the field of algorithm animation this is a commonly advocated policy, Brown (1984). In CAD systems the user can frequently view the data–structure in several different ways, either in the literal sense as a wire frame, or, as a shaded, or, ray-traced model. The object may also be viewed as a list of component parts, or as a library of user defined primitives. With WET & STICKY the user could view the standard colour image plus colour coded maps showing the distribution and values of individual attributes. In most operations these alternative views need not be exactly the same size as the image to be effective. However if they were the same scale then they would offer the advantage of allowing the user to accurately paint or re-set an attribute in its own attribute map window.

An additional aid suggested by Waite (1990) would be that of user defined areas within the image. For example, the user could use a rubber-band selection method to select some arbitrary region of the image and give it a text label, then he could use this as a descriptor of this area. Alternatively the region could be defined by painting an area. If the user wanted to carry out a global operation to this region of the image he would then call up a dialogue box which would display all the attributes available and then reset explicitly those belonging to the canvas or paint within that area. Within the dialogue box the user could also choose from a menu any of the other areas which he may have named.

This system would make the notion of accessing and re-accessing logical areas within the image such as the sky, or foreground, for example, much easier. In many ways this approach is a combination of the easy unit editing of object-oriented drawing systems with the discrete editing associated with pixel based systems.

There already exist systems that go some of the way along this path. Canvas (1988) the Macintosh ™ paint/drawing system supports two types of image layers. One is a painting layer, and the other a object oriented drawing layer, which itself may be multi-layered. An area from the paint layer may be transfered to the drawing layer in the process becoming a drawing object, and vise–versa. Canvas does not allow for the naming of drawn objects although many other systems do support this feature, (MacDraw 1990). Others such as PixelPaint Professional (1990) make use of colour masks, where the user can define which colours should act as a mask.

Waite also proposes that it is not the element of the image that is defined and named like a drawn object , rather that an area of the canvas is defined independently of what is upon it at the time of definition. Some systems allow the user to define a mask which is used to protect an area of the image. But in most cases only one, the most recent mask is stored, rather than several. In addition there is usually no provision for applying global operations to the area included or excluded from the mask. The mask is used as if it were a physical device which obscures and protects an area of the image from paint applied to the image.

For greater flexibility we should extend Waite's proposal so that the user should also be able to edit these named areas in order so that they may match changes to the overall composition of the image.

8.4.2 Three dimensional viewing

As we have seen each canvas cell has knowledge of the amount and characteristics of any paint it contains and knowledge of its own attributes. The standard way of viewing the image would be in two dimensions, with only the colour of the paint or canvas being used for display. The marjority of the information held by the canvas cells is not used for display but is used instead by the painting engine to model the flowing and mixing of the paint.

However, each cell contains sufficient information for it to be displayed in three dimensions. The only additional information required is one or more light sources. The volume of paint contained by the cell can be equated to its height, and the liquid content can be used as a measure of the paint's glossiness. Combined with the paint's colour one can use this information to make a bump map of the image.

The painting engine is responsible for the updating of the display. At any instant it is only ever working on a very small part of the image, that is the canvas cell which it is processing and its immediate neighbours. This would make it only a relatively small additional task to update a bump map locally. Given sufficient processing power it would be possible to interact with the bump map and to paint effectively using three dimensional paint. For this to work satisfactorily the system would also have to employ complex brushes as described in the further work section. A bump map version of WET & STICKY has not yet been implemented as this is not considered to be within the scope of this thesis.

8.5 Summary

In this chapter I have proposed additions to the basic model, and more powerful derivations of the basic model. The additions would in the case of, user–definable areas, improve the useability of the model, and in the case of three dimensional viewing, improve the appearance of images produced on with the model.

Each derivation gives the user a greater degree of control over the image and how the paints will evolve. A configurable system would allow exact editing of attributes. User definable rules allows for experimentation with alternative realitites. It is impossible to determine if the advancements proposed by the data configurable and rule configurable systems would be truly advantageous, liberating, confusing, or beneficial to the artist. They are described to show the power, potential and flexibility of the model. No qualitive pre-judgements are made.

Appendix A

Paint Smith (1982)

This appendix will describe the basic functions of Paint Smith (1982). This description will serve as a basic list of the type and range of functions available to most subsequent paint systems, i.e.;

- Basic drawing operations such as, circle, rectangle, irregular shapes. These could be drawn with or without fills.
- The user could design his own brush, by painting the shape of the footprint of the new brush. This can be saved with a user defined name and could subsequently be selected from a menu along with the standard brushes provided by the system.
- Saving images to files and retrieving images from files.
- While working on an image the user could use this system to cut and paste portions of the current or any other image.
- Selecting colours from a palette of available colours. The user could, customise an existing palette or create his own. This was done by selecting a colour from the palette and adjusting its hue, saturation, and value using virtual sliders.
- Magnification of the image up to eight times.
- The ability to draw with antialiased lines of variable thickness. It is interesting to note that on the menu this function was called sketch. To an artist to sketch infers a rougher, more unfinished technique than normal rather than, as in this case the reverse, a smooth line.
- Record and play-back a set of actions.

- Specialized paint functions such as, smear, slide, z-paint and z-picture paint. Smear and slide are self explanatory#. Z-paint gives the effect of the paint on the brush going over some paints and under others. This is achieved by only displaying pixels whose colour value is lower than that on the brush. Z-picture paint is the inverse of z-paint with only pixels whose colour is higher than that on the brush being displayed. The picture element refers to fact that the brush used with this paint is itself an image.

Recently many new additional effects have been added to paint systems, some are spatial distortions, such as the ability to give the effect of mapping the image onto a sphere, others allow the user to apply mathematical filters to areas of the image, and some also provide tools for altering the brightness, contrast, and colour of the image.

Smear averages the colour of pixels under the brush at the start of the stroke with those along the stroke. Slide shifts the colours under the brush at the start of the stroke along the direction of the stroke. In both cases the effect is limited to a short length from the start of the stroke. Sometimes this distance is user definable.

Appendix B

Hairy Brushes Strassman (1986)

Strassman abstracted the process of painting a sumi-e brush-stroke into three basic components whose individual behaviour and characteristics interact to created the final image.

The bristles of the brush.

The brush is considered as a collection of bristles. Each has its own supply of ink and a position relative to the brush handle. In complex brushes the bristles may move relative to one another, i.e. when pressure is applied. As the brush is moved along a trajectory calculations are performed which:

- Update the state of the bristle in accordance with one or more rules applied to the colour, ink quantity, relative position and other bristle properties.
- Transfer an image of the bristles onto the paper. (In a typical case each bristle would contribute one pixel. Ink is transferred only if the pressure is sufficient and there is ink remaining in the bristle).

The dip of the bristles into the ink.

The complexity of traditional sumi-e brush-strokes is a factor of the amount and distribution of ink amongst the bristles. By separating control of the dip from the brush and stroke the user can use the same brush with different dips to achieve different effects.

The trajectory and pressure of the stroke.

A trajectory is simply the position and pressure of the brush through a period of time.

Using a catalogue of brush-stroke rules, and allowing the user to make decisions about which rules are to be applied, how they are implemented, and how the components interact, he was able to model many different types of *bokkotsu* strokes.

The user could select the amount of ink held by the brush. The distribution of the ink amongst the brushes bristles was also controllable as was the behaviour of the ink between those bristles with and without ink. Another variable was how the bristles behave when pressure is applied to the brush. Either more bristles come into contact with the paper, or all the bristles are in contact all the time and spread under the influence of pressure, or a combination of the two.

Because the system was not real-time it allowed editing of the trajectory before it was painted and it also offered the option of storing the trajectory as a separate element. This enabled the user to play around with the ink, dip and bristles parameters and apply different combinations to the same trajectory. Knowing the length of the stroke also allowed the user to preset the portion of the stroke that would be in solid ink as opposed to feathered or streaked ink. This made a nonsense of allowing the user to dip the brush since the choice of the amount of ink on the brush was overridden after the stroke.

In most paint systems a two–dimensional array of points, representing the footprint, is moved through a trajectory a bit at a time. After each move a copy of the brush pattern is copied onto the frame buffer. To anti-alias the image a super-sampled patch is used. The brush is moved across a virtual screen, 4 or 16 times the resolution of the screen. This trajectory on the virtual screen is then reduced to the resolution of the display screen.

Strassman found several problems with the 2D brush. Because of hardware limitations (a mouse and function keys as input) it was very difficult to specify rotation. When using a real brush the bristles flow behind the shaft of the brush and rotate with the rotation of the shaft. A second problem was bristle spreading . If the bristles are arranged in a square lattice what happens when pressure is applied? What happens to the gaps between the bristles? He also argues that a 2D brush gives a predictable start and finish to the stroke, thus "wasting" the effects generated in the middle.

Of these the bristle spreading seems to be the hardest problem to overcome. A stylus with rotation sensors would solve the first problem, though one could disagree with his view on the predictability of the start and finish of the stroke. This is a function of the ink flow model as well as the bristles footprint. With real brushes the application of pressure, by increasing the spaces between bristles and thus decreasing the ability of the brush to *hold* paint, induces the flow of paint. The way paint flows differs amongst the bristles, thus any pressure applied at any time throughout a stroke will produce a slightly unpredictable result. The degree of un–predictability could be variable, but the result would be non–uniform starts and finishes of strokes.

His second implementation used a 1D brush. The path was represented as a series of nodes connected by line segments and the pressure gave the width at each node. A trapezoid was first drawn over each line segment and a transition elbow at each node. The difference in width between nodes was linearly interpolated. The problem of rotation is removed because the width is always perpendicular to the path. Spreading is now defined by the width which is a function of the pressure applied However there were still some problems which led to a third implementation. The main problem was that the elbow pixels were often missed. They were also missed when the width changed.

The third implementation used polygons and an algorithm which when given a polygon with V vertices , and a N -dimensional vector value at each vertex, which generates;

•A list of all the pixels contained by that polygon.

•For each pixel, an N-dimensional vector which is the linear interpolation of the values at each of the vertices.

Now each of the pixels is drawn only once. They are drawn in chronological order, and the missing pixel problem is eliminated. He describes the main advantage of this system as a new way of looking at the problem. The question he was originally asking was, here's the brush, what pixels does it paint onto? He turned it on its head. Now the question is, here are the pixels, from what part of the brush did they get drawn ?

Complex brushes for the WET & STICKY model

The use of hairy brushes with an early implementation of the WET & STICKY model showed that the effect of wet paint was to almost instantly remove any of the subtle marks they produced. The requirements of a complex brush which can take full advantage of the WET & STICKY model are outlined below.

Brushes are extremely versatile and complex tools. Within the mark making tools provided with the average computer based painting systems the brush is only one element, with the marjority of tools designed for the drawing of regular shapes. In the studio it is the brush which is the most important tool. The artist will usually have a vast array of different types and sizes of brushes.

The basic attributes which can be used to model and classify brushes are,

the number of bristles,

the profile of the brush, that is the varying cross-section of the brush at different heights,

the flexibility of the bristles,

the ability of the bristles to hold paint, the distribution of paint through the bristles,

the behaviour of the bristles under the application of pressure, or acceleration/deceleration during a stroke.

Hairy brushes only considers, the number and flexibility of bristles, their ability to hold paint and how they spread under pressure. Furthermore it only considers these attributes in a limited manner. This is primarily because hairy brushes were designed to work with the Shoup model. Once wet paint is introduced then one must additionally model,

a brush with stiff bristles, when pushed hard, will plough through a wet substrate paint ,

a soft brush loaded with liquid paint will glide across the surface of the substrate paint,

a fully loaded soft brush may splatter paint when subject to a sudden change in acceleration or deceleration

a dry brush may pick-up substrate paint.

All these affects, and many others, are the result of a combination of the qualities of the brush, the paint on the brush, and the attributes of the substrate and, or substrate paint. To adequately model brush dynamics would require a very complex model. Such a model could be constructed using the same stage-set methodology emploied in WET & STICKY. A further problem is how to handle the processing of the application of paint. The amount of processing required suggests that it would be very hard to achieve a real time modeling of brush dynamics. Even existing systems based on the Shoup model can have difficulty in tracking the brush and updating the screen.

The problem of building a adequate brush model is on parr to that of building the WET & STICKY model. It is for this reason that, although it is a very attractive research area, WET & STICKY does not provide any support for complex brushes.

References

Adobe (1985)	Adobe Systems Incorporated *"PostScript Language Reference Manual"* Addison Wesly Redding Mass. 1985
Adobe (1990)	Adobe Photoshop™ 1.0.7 Adobe Systems Inc., 1989–90
Aldus (1990)	Freehand Application Version 2.0, © Altsys Corporation, 1988
Ashton (1972)	Dore Ashton *"Picasso on Art A Selection of Views"* Thames and Hudson p 72
Beach etal (1982)	Richard J. Beach, John C. Beatty, Kellogg S. Booth, Darlene A. Plebon *"The Message is the Medium: Multiprocess Structuring of an Interactive Paint Program"* Computer Graphics Vol 16, Number 3 July 1982
Berger (1963)	John Berger *"Selected Essays and Articles, The Look of Things"* Penguin Books p 165
Blesser et al (1988)	T. W. Blesser , J. L. Sibert, P. J. Magee *"Charcoal Sketching: Returning Control to the Artist."* ACM Transactions on Graphics Vol 7 No 1 January 1988.
Blinn (1978)	J Blinn, *"Simulation of Wrinkled Surfaces"*, Computer Graphics (Proc. SIGGRAPH), Vol. 12, No. 3, August 1978, p 286 – 296
Blunt (1960)	Anthony Blunt *"Artistic Theory in Italy, 1450 – 1600"* Oxford Paperback 1960.

Brown (1984) — Marc H. Brown, Robert Sedgewick
"A System for Algorithm Animation"
Computer Graphics, vol18, no. 3, July 1984
p 177

Canfield et al (1982) — Canfield, David Smith, Charles Irby, Ralph Kimball, Bill Verplank, Eric Harslem,
"Designing the Star User Interface"
Byte Magazine, vol 7, no. 4, April 1982.

Canvas (1988) — Canvas 2.0
©Deneba Software, Inc
1988

Carroll at al (1980) — John M. Carroll, John C. Thomas
"Metaphor and the Cognitive Representation of Computing Systems"
Research Report, Computer Science Department,
IBM Watson Research Centre, Yorktown Heights.
1980

Cennin (1954) — Cennino d'Andrea Cennini
The Craftsman's Handbook " Il Libro dell Arte" .
Translated by Daniel V. Thompson, Jr.
Dover.
1954

Chipp(1968) — Herschel B. Chipp
"Theories of Modern Art"
University of California Press
1968
p 273

Chipp (1968) — Herschel B. Chipp
"Theories of Modern Art"
University of California Pres
1968
p 548

Cockshott (1989) — Tunde Cockshott
"Report on the current work at the graphics and animation labroraties at the Media Lab, M.I.T. Boston."
Glasgow University Computing Science
Departmental Research Report 1989
Glasgow University Department of Computing Science.

Codd (1968) — E. F. Codd
"Cellular Automata"
New York: Academic Press Inc
1968

Collins et al (1987) — Collins J., Welchman J., Chandler D., Anfam D.,
"Techniques of Modern Artists"
Tiger Books International, Quarto Publishing London
1987

Cook (1984)	Robert L. Cook *"Shade Trees"* Siggraph Vol. 18, No 3. July 1984
Donaghue et al (1990)	Donaghue, Karen, and Knowlton, Ken *"Computer Simulation of Calligraphic Pans and Brushes."* Internal report, Freestyle/Advanced Technologies Group, Wang Laboratories, Inc 1990
Em (1983)	David Em *"Butterfly Nets Revisited: The artist in the Lab"* ACM Siggraph 1983 Vol. 18
England (1987)	David England *"A Sun User's Guide"* Macmillan Education 1987
Foley and van Dam (1990)	James Foley, Andries van Dam *"Computer Graphics Principles and Practice"*, Second Edition Addison-Wesley Publishing Company 1990 p 584 – 598
Giloth & Veeder (1985)	Copper Giloth and Jane Veeder *"The Paint Problem"* IEEE Computer Graphics and Applications July 1985, p 67 – 75
Goldberg and Robson (1983)	Adelade Goldberg and David Robson *"Smalltalk-80: The Language and its Implementation"* Addison-Wesley, 1983
Gombrich (1978)	*"The Story of Art"* E.H. Gombrich Phaidon Press Limited 1978 p 438
Grandis (1986)	Luigina de Grandis, *"Theory and Use of Colour"* Blandford Press, 1986
GraphistPaint (1988)	GraphistPaint©II–Colour by E. Wenger & S. Rostan Degrader option Adone Systeme SA 1988
Greene (1985)	Richard Greene *"The Drawing Prism - A Versatile Graphical Input Device"* Siggraph 85 Vol 19 No 3, p 103 - 110

Greiman (1990) April Greiman
"Hybrid Imagery: The Fusion of Technology and Graphic Design"
Architecture Design and Tecnology Press 1990.

Haeberli (1990) Paul Haeberli,
"Painting by Numbers: Abstract Image Representation."
Computer Graphics, Volume 24, Number 4,
August 1990.

Hammersley & Handscomb (1964)
J M Hammersley and D C Handscomb,
"Monte Carlo Methods",
Methuen & Co. Ltd. 1964

Hartson (1990) Hartson R. H., Gray P. D.
"Temporal Aspects of Tasks in the User Action Notation"
Journal of HCI, to appear.

Hayes (1978) Hayes P. J.
"The Naive Physics manifesto"
in Expert systems in the microelectronic age.
Editor D. Mechie
Edinburgh University Press.
1978

Lewsi (1983) John Peter Lewis
"Texture Synthesis for Digital Painting"
Computer Graphics Volume 18, Number 3
July 1983

MacDraw (1989) Macdraw II 1.1v2® Claris™ Corporation
April 1989

Norman (1986) Donald A. Norman
"Cognitive Engineering"
User Centered System Design
Editors Donald A. Norman, Stephen W. Draper
Lawreence Erlbaum Associates Publishers
p 47

Oasis (1990) Oasis Image Creation Software
Time Arts
Santa Rosa, California
1990

Oddy & Willis (1991) Robert J. Oddy and Philip J. Willis
"A Physically Based Colour Model"
Computer Graphics Forum, Volume 10, Number 2,

Owen (1986) Owen D.
"Naive Theories of Computation"
page 187 in User Centered System Design
Editors Donald E. Norman, Stephen W. Draper
Lawrence Erlbaum Associates Publishers
1986

Patterson (1990) — John Patterson
personal communication,
Computing Science Department, University of Glasgow,
1990

PixelPaint Professional (1990)
PixelPaint Professional™ 1.0,
©Pixel Resources, Inc.
1990

Reeves (1983) — W T Reeves
"Particle systems – a Technique for Modelling a Class of Fuzzy Objects"
ACM Transactions on Graphics
Vol. 2, No. 2, ACM, 1983

Richet (1987) — Catalogue by Michele Richet
"II Drawings Watercolours Gouaches Pastels"
The Musee Picasso, Paris,
Thames and Hudson.

Richter (1965) — Hans Richter
"Dada Art and anti Art"
Thames and Hudson
p 51

Rivlin (1986) — Robert Rivlin
"The Algoithmic Image"
Microsoft Press 1986

Schiff (1986) — Catalogue by Gert Schiff
"Picasso at Home, Selections from the Marina Picasso Collection",
Center for the Fine Arts, Miami, Florida,
Nov 85 – March 86.
p 100

Schnabel (1982) — Julian Schnabel,
"Catalogue for the Tate Gallery Exhibition"
1982, p 6, "Death of Fashion"

Shepard (1989) — Andrew Shepard,
"Analysis and Training in Information Technology Tasks Task Analaysis for Human Computer Interaction",
Editor Diaper
Ellis Horwood Limited 1989

Shneiderman (1987) — B Shniderman,
"Designing the User Interface: Strategies for Effective Human-Computer Interaction",
Addison Wesley,
1987

Smith (1974) — Edward Lucie Smith,
"LateModern: The Visual Arts Since 1945"
Praeger
1974

Smith (1982)	A. R. Smith *"Paint"* in Tutorial: Computer Graphics, (J C Beatty, K S Booth eds.) IEEE Computer Society Press 1982 p 501 – 515
Smith (1987)	Randall B. Smith *"Experiences with the Alternate Reality Kit: An example of the Tension Between Literalism and Magic"* IEEE Computer Graphics andApplication Sep 1987 p 42 – 50
Smith (1978)	A. R. Smith Paint, Technical Memo. No 7, Computer Graphics Lab New York Institute of Technology, 20 July 1978
Smith (1989)	Reffin Smith *"Beyond Computer Art"* from Computer Art in Context Panel session ACM Siggraph 1989
Sokolove (1990)	Deborah Sokalove *"The Image in the Magic Box"* Art Journal Fall 1990. p 272
Strassmann (1986)	Steve Srassman *"Hairy Brushes in Computer–Generated Images"* MSc Thesis, MIT, June 1986
Sutherland (1963)	Sutherland I. E. *"Sketchpad, A Man – Machine Graphical Communication System"* PhD thesis M.I.T. 1963
Vermeulen (1989)	Allen H. Vermeulen and Peter P. Tanner *"PencilSketch – A Pencil Based Paint System"* Graphics Interface 1989 p 138
Waddington (1969)	C. H. Waddington *"Behind Appearance : A Study of the Relations between Painting and the Natural sSciences in this Century."* Edinburgh University Press 1969. p 158
Waite (1990)	Kevin Waite personal communication, Computing Science Department, University of Glasgow, 1990
Walker (1983)	John A Walker *"Art in the Age of the Mass Media"* Pluto Press 1983 p 40

Ware & Baxter (1989)
Colin Ware and Curtis Baxter
"Bat Brushes: On the uses of six Position and Orientation Parameters in a Paint Program"
ACM CHI 1989 proceedings
May 1989

Watters et al (1987)
Geoff Watters, and Philip Willis
"Ultrapaint : A new Approach to a Painting System"
Computer Graphics Forum 6 1987 pages 125 - 132

Whitted (1983)
Whitted , Turner,
"Anti–aliased Line Drawing Using Brush Extrusion"
Computer Graphics 17, 3, July 1983,
p 151–156

Willis (1984)
P. J. Willis
"A Paint Program for the Graphic Arts in Printing."
Eurographics 1984 Conference Proceedings
Elsevier Science Publishers B. V. North Holland
p 109 – 120

TITLE WET AND STICKY: A NOVEL MODEL FOR COMPUTER-BASED PAINTING

AUTHOR Malcolm Tunde COCKSHOTT

DEGREE Ph.D

AWARDING BODY Glasgow University

DATE 1991

THESIS NUMBER DX205420