ADVISORY GROUP ON Computer Graphics

The DESIGN of Virtual Environments with particular reference to VRML

Principal author **Stephen Boyd Davis**
Contributors Avon Huxor, John Lansdown
Centre for Electronic Arts
Middlesex University June 1996

sima
Support Initiative for Multimedia Applications
Part of the JISC New Technologies Initiative

Contents

About this report

This report was written by Stephen Boyd Davis, Principal Lecturer in Multimedia at the Centre for Electronic Arts, with contributions by colleagues in the Centre.

Sections on Fundamentals of spatial organisation and Movement Perception, together with the paragraphs on Architecture as spatial design were written by Prof John Lansdown, Emeritus Professor. They are derived from a report on Visual Perception prepared for the Human Factors Unit, British Telecom Laboratories, Martlesham, and are used here by kind permission.

The sections on designing within current and future versions of VRML were written by Dr Avon Huxor, post-doctoral researcher in virtual environments, who also contributed sections on interaction in social spaces and augmented realities.

Acknowledgements

The authors would like to thank the following:
Prof Roy Armes
Tracey Durkin
Charlie Gere
Nat Goodden
Dr Huw Jones
Andy Wilson

Helena Athoussaki and Grethe Mitchell, postgraduate students of the MA Design for Interactive Media, built a VRML model as the other component of this SIMA Project.

We thank BT for supporting the development of the Centre's 3D-Workshop which made the student project possible, and particularly Andrew McGrath for his technical and tutorial assistance for student projects.

The report was funded by a grant from the Advisory Group on Computer Graphics: Support Initiative for Multimedia Applications.

Introduction

Two key points underpin this report. The first is that the construction of Virtual Realities should be seen as an intentional activity, based on thoughtful well-informed and inventive decision-making. It is a design process in the best and broadest sense of the term.

The second is that the decisions in that design process should ultimately be based on effectiveness. All the decisions made can be evaluated by asking: will this project, in this form, achieve the desired result?

Definition

For the purposes of this report, we consider Virtual Reality to be

> *a computer-based representation of a space in which users can move their viewpoint freely in real time*

Worlds created in VRML-1 fulfil this definition (with a caveat about real-time movement in most cases).

Fuller definitions of VR are possible, and desirable. For example, many would expect users to have the ability to interact with objects in the space, not just to alter their view of it. Others would expect the objects to be capable of autonomous animation, without the user's intervention. VRML-2 provides these things. To avoid being tied to any one implementation, this report largely steps back from these contingent details and tries to address fundamental issues.

Some believe VR should be limited to those environments which mimic external realities but there are problems with this approach:

- Mimicry means deciding what we mean by reality (this is not just a philosophical nicety).
- Mimicry excludes from the discussion lots of potentially interesting applications, for example 3D models of numbers ('data visualisation').
- Complete mimicry is impossible—some process of selection, abstraction and representation is inevitably involved.

The challenges of VR are not only technical

There has been a tendency for VR to be hailed as an unproblematic answer to many problems in computing, from database, through educational technology, to virtual museums. The task of making virtual environments is seen as a technical one—essentially of reconstructing reality in the computer. Indeed this is seen by many as a sufficient definition of VR. However, there are clear problems.

Simulating reality is not enough

If we take as an example the idea of a virtual museum, are we to limit ourselves to placing objects in fixed locations in a building of a given appearance? What is the point? It could be argued that a straightforward visual database is a better 'virtual museum' since it allows works to be located within as many dimensions as there are fields in the record. What do we gain by making a virtual physical construction? Nothing, unless we abandon some of the simplistic assumptions about virtuality. We are faced with a rich variety of choices. What should inform those choices? Even supposing that the replication of reality were possible in VR, there would be many occasions when there was little to gain by doing so.

The aim of this report is to highlight those areas of promise and concern in relation to VR which have not been dealt with by others.

Virtual Reality requires design

It is possible to deal separately with some key characteristics of VR:

1 Spatiality—the uses of space.

2 Virtuality—directness, the feeling of 'being there'.

By this distinction we may be able to usefully disentangle meaning within a space (spatiality) from the actual relationship of the user to that space (virtuality).

Connecting the two themes is the issue of...

3 Representation

Even when we have decided what to model, there are still many choices available to us about how to model it. A representation is not the thing it represents. Even the most ambitious virtual reality will not be reality. It is clear that we are dealing with a representation: as creators of such environments we have both the obligations and the freedoms that any designer has in constructing a representation.

A note on the five senses

Virtual Reality almost always privileges vision. Sound may also be used, simply or in more sophisticated ways (for example using phasing). Where additional hardware is available, tactile feedback and the sense of body orientation may be possible. Smell and taste are currently rarely addressed. In this it follows the limitations of cinema.

The construction of meaning

It might be argued that a virtual reality is no more a representation of something else than a building is a representation of something else—it is simply a construction. Whether it is regarded as representation or construction, however, there should be no doubt that it has the ability to carry meaning, in the same way that a building can present the intentions of those who built it. There are many other constructions which also carry meaning, including theatre design, interface design and product design. In addition, interesting constructions are described and alluded to in various forms of literature, and some of these offer interesting ideas for the design of VR. This report attempts to draw out the lessons of these various approaches to designing with space.

VRML

VRML—the Virtual Reality Modelling Language, often pronounced 'vermal'—is the standard being developed for putting virtual realities onto the Internet. For many in Higher Education it will be with this set of technologies, rather than any other, that they will experience virtual reality. This report concludes with a snapshot of VRML now, some ideas about its future, and recommendations on how the whole concept of VR can best be harnessed for our needs.

We hope that this report helps to stimulate a wider debate about the design and use of VR. We would also like to think that it is useful in part as tutorial material for students, whose engagement with these issues should be actively encouraged.

The uses of spatiality

What can we use space to mean? What do we use space for in real life? Space can be physical; it can also be metaphorical. We should not omit the uses of 2D space, which may also help us in designing in 3D.

For example:

- Height is used to indicate ‘more’ of something (including metaphorically even when no real height or quantity is involved, such as a high temperature or a high note).
- Proximity and alignment in 2D are used in computer interfaces to show which options belong together.
- Thrones are higher than commoners’ seats to show social power.
- People in social situations use proximity to show how intimate they want to be with others in the space. Orientation is also used, facing someone or turning away from them.
- Large architectural spaces are generally felt to invoke awe, while small ones may be oppressive or cosy, depending on a number of factors.
- Using different length lenses (telephoto, wide-angle etc) creates different kinds of perspective: the way scenes are shown alters the viewer’s perception.

The use of fully dimensional space is one of the defining characteristics of VR. So far discussion of three-dimensional objects and worlds in VR has tended to look at the technology—what is possible—and the application—what is demanded. However there is an entire layer missing from the discussion, a layer between the technology and the application, which should deal with the qualities of spatiality itself. In other words: ‘what can we use Space for?’

We can begin to answer this question by looking at how space has been used in other media and other contexts, rather than just within VR. As with many applications of computing, there is benefit in intelligently applying critical ideas from other media and other technologies, rather than looking only in our own backyard.

Do we know what space is?

Readers might be impatient at being asked what space is. After all, does it matter? Must we waste time on definitions? We suggest that some awareness of the debate about definition is needed, especially in the context of transcultural software development.

> It may be that all human beings have the same perception of space at the biological level of perception. But certainly every society uses its space differently, both technologically and artistically.
>
> J David Bolter 1986 Turing’s Man p80

Bolter’s may be is important. We should look briefly at the question of biological perception, before considering his second point.

The two opposing schools of thought can be briefly characterised as follows:

- The perception of space is a culturally influenced phenomenon
- Space is perceived identically, as a universal human phenomenon

This is the nature/nurture controversy in one of its many guises: do we perceive space according to universal optical and perceptual principles on which social and environmental conditioning has no effect; or do we acquire some aspects of our perceptual system from external influences? It should be emphasised that we are leaving aside for now the question of its representation. If the very perception of space is culturally determined, then a new set of variables enters the equation.

The first widely acknowledged suggestion that the perception of space might be culturally determined arises in the work of Edward Sapir and Benjamin Whorf. Their work and that of their followers has acquired a remarkable popularity, becoming in some circles an orthodoxy in its own right. A favourable view of Whorf's work is found in a paper prepared for AGOCG (Hopgood 1993) which was prompted by a debate on standards for specifying time and space in the PREMO standard.

The key argument of the Sapir-Whorf hypothesis is that perception of space is determined by culture, and particularly by language. Whorf (1897-1941) was amongst other things an amateur linguist, who studied several languages including Hebrew, Aztec, Mayan, Hopi and Chinese. His view (in the words of Hopgood) is that:

> Europeans have a notion of time and space that is generally assumed by them to be universal. This gratuitous assumption is naive, arrogant and wrong.
>
> Hopgood 1993, p3

Hopgood is confident that the Hopi language contains no reference to time either explicit or implicit (ibid, p3). His argument in relating Whorf's theories to the development of standards for multimedia, is that, since no language system (or concept-system, in the view of this school of thought the same thing) can claim to apprehend space or time with greater truth than another, our reason for preferring one to another should be based on functional grounds. We should choose a model of space-time to serve a purpose, rather than trying to assert that it is somehow true. The criterion should be: does it work? For example, Newtonian physics is useful for prediction of everyday cause and effect, but unhelpful under special conditions such as high speeds or sub-atomic scales. Under these conditions we have to pick another model.

Transformative technologies

Most people, even where they disagree with the Sapir-Whorf hypothesis that language and culture influence perception itself, would endorse the idea that forms of communication are transformative. In other words, what can be said and thought is influenced by the means selected (or available) for expressing or representing it. The important consequence would be a recognition that Virtual Realities will themselves facilitate the communication of some kinds of information and experience, and militate against others.

To the Victorians, the idea that media are not transparent to the ideas they convey would have been unfamiliar. The whole period since the Renaissance has seen a largely unspoken belief that media are an unproblematic window on reality. This belief has been largely superseded during the twentieth century by that of transformative technologies. Almost notorious is Marshall McLuhan's dictum, 'the medium is the message'. Postman (1982, 1985) has made extensive use of a McLuhanite approach in his criticism of television and other predominantly visual media as against the merits of the word. Ong devotes all of Orality and Literacy (1982, p155) to the transformative effect of writing and printing on the world of ideas, and argues explicitly that 'These technologies...style what we know in ways which make it quite inaccessible and indeed *unthinkable* in an oral culture,' (our emphasis).

Pinker (1994) aligns himself strongly with the Nature camp against the Nurture-based view, and adamantly refutes the Whorfian hypothesis. He does so both on the grounds that Whorf and his followers gathered faulty evidence, and that they argued falsely from these premises. It may be that there is no possibility of deciding between the opposed theoretical positions. However, in the context of harnessing virtual realities for particular purposes, it would seem that the functional view advocated by Hopgood is a productive one. The implication is, that we should represent in our virtual realities only those aspects of reality which serve our purpose at the time. This would suppose at the very least a selectivity in the attributes of reality which we might describe, and perhaps the deliberate construction of attributes which are 'unrealistic' but useful. It would also imply Virtual Object and Environment data which can be differently filtered and represented on an as-needed basis, rather than which have a fixed form.

Fundamentals of spatial organisation

The process of perception is not of merely scientific interest. It is regrettable that so many visual artefacts of all kinds are created in ignorance of important theories about perception and cognition: many designs are based on inherited, ill-founded beliefs. However there are also many traditional practices in design which do seem to reflect well-founded principles of spatial organisation. These are broadly consonant with the approach of Gestalt psychology, and include the following important principles of perceptual organisation—anyone embarking on the design of spatial artefacts should be aware of these possibilities and constraints. Some of these findings seem like 'common sense' (for example in relation to Proximity) while others are surprising (Orientation), or unfamiliar (Symmetry).

Smooth continuation

We group together in a single structure those parts which seem to align or continue smoothly. Thus, in Figure 1, we see two curved lines crossing at right angles (as in Figure 2) rather than two V-shaped forms meeting at a point as in Figure 3. This principle is also known in some texts as 'good continuation'.

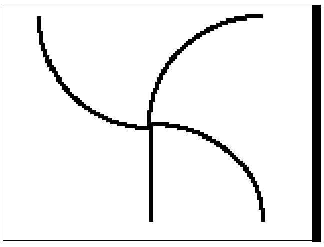

Fig 1 Smooth continuation of lines

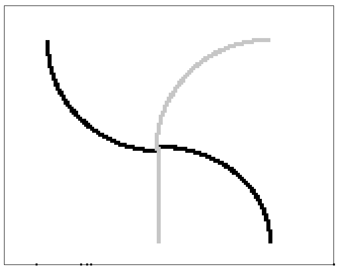

Fig 2 Preferred visual grouping of lines

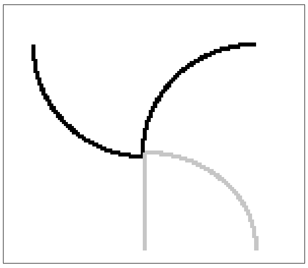

Fig 3 This visual grouping of lines is normally rejected (even though it was the grouping that was used to create Fig 1)

Perhaps this feature derives from our experience of the 'objectness' of things. Thus, we see objects as a whole even when they are partially occluded by other things because we know that they do not break up when they disappear from view, as Figure 4 shows. Objects overlaying others have continuous outlines. The overlaid objects have interrupted outlines but, by the principle of smooth continuation, we can deduce continuity behind the overlaying object. Kellman and Spelke (1983), however, show that 3-4 month old infants only appreciate this if the objects

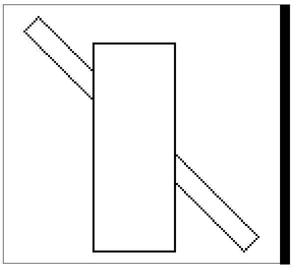

are moving relative to one another.

Fig 4 Smooth continuation contributes to our feeling of 'objectness' and allows us to judge that the thin rectangle might be a complete object behind a rectangle in front

Presumably we learn to judge the still-image situation from experience although, as Figure 5

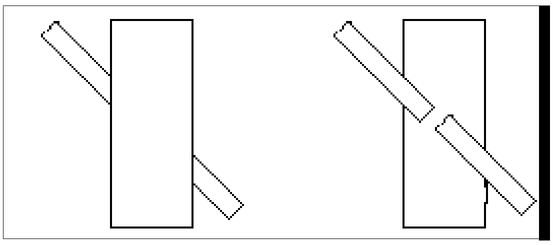

illustrates, our tendency to accept objectness through the principle of smooth continuation can sometimes mislead us.

Fig 5 Sometimes objectness overrides smooth continuation and gives us the mistaken impression that lines continue when they do not. Not only is the diagonal rectangle in the left hand side of the figure not continuous, the two diagonal rectangles (as can be seen in the figure to the right) do not even align

Proximity
We group together those parts that are closest together. As can be seen in Figure 6, we perceive the group (a) as three vertical lines of dots and the group (b) as three horizontal lines of dots. The dots in (c) are equally spaced and do not suggest an orientation.

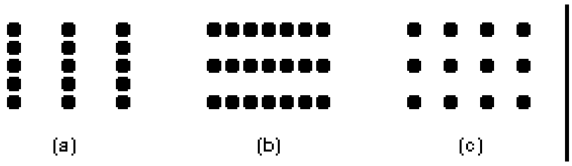

Fig 6 The principle of proximity determines our interpretation of the groups

Obviously proximity and the size of the elements that make up the pattern are related factors here (Zucker and Davis 1988).

Similarity
We group together those parts that appear 'similar'. Hence in Figure 7, we see separate white diagonal lines and black diagonal lines rather than vertical or horizontal lines of black and white dots.

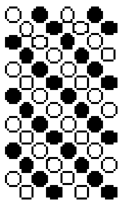

Fig 7 The principle of similarity determines our interpretation of the groups in this case and thus we tend to see diagonal lines

Sometimes, similarity can override proximity as the organising principle (Figure 8).

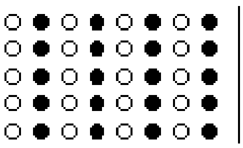

Fig 8 The principle of similarity ensures that we group the dots in vertical columns

Orientation
We group together items which are arranged in a vertical or horizontal orientation in preference to those orientated on different axes. Thus orientation seems often to be a stronger grouping principle than similarity (Figure 9).

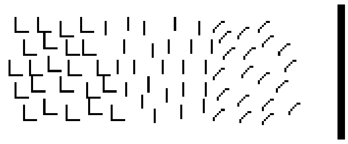

Fig 9 We group together items having similar orientations rather than similar shapes

Orientation, as we shall see later when we look at relationship to frame, also affects our interpretation of a shape.

Closure

We group together parts that give the appearance of closed shapes (Figure 10).

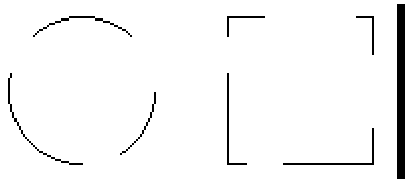

Fig 10 The principle of closure determines that we see these interrupted lines as forming closed figures

Thus it is often possible to suggest a virtual frame around a figure by only drawing its corners. The organisational principle of closure seems to come to the fore when we interpret sketch drawings—which are often incomplete but which we normally have little difficulty in understanding (Figure 11). See also Productive ambiguities in drawing p66.

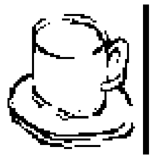

Fig 11 Closure comes into play in our recognition of sketches

Relative size: figure and ground

Given two superimposed areas, we will tend to see the smaller as a figure against the larger background rather than vice versa (Figure 12).

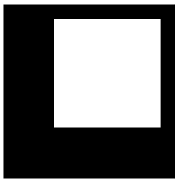

Fig 12 We tend to read this image as a white square on a black one rather than a black square with a hole in it

When there is little difference in the size of the parts, ambiguity can result and we are unable to fix exactly which is the figure and which is the ground. Sometimes this ambiguity can be exploited for art purposes. Figures 13 illustrate this.

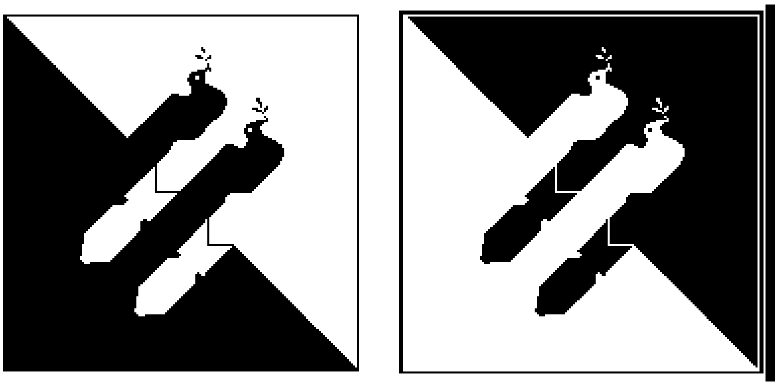

Fig 13 Two versions of an image in which the relative sizes of the parts are similar so that it is difficult to distinguish figure from ground (based on a CND poster)

In cases like these we can chose arbitrarily which is the figure and which is the ground. One choice allows us to perceive one meaning, the alternative choice allows us to perceive a different meaning. This effect is the basis of a number of well-known optical illusions such as the faces/vase illusion of Figure 14.

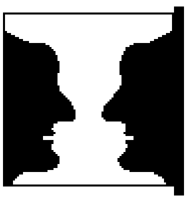

Fig 14 Different information arises depending on whether we take the black or the white to be the ground

Part of the ambiguity of Figure 14 seems to derive from the mixture of convexity and concavity of the black and white parts. Kanizsa and Gerbino (1976) show that shapes that are symmetrical about a vertical axis are usually seen as figures against ground but that this is not always the case if the forms are concave. Carrying out experiments with diagrams similar to those shown in Figure 15, they discovered that over 92% of people tested see the convex shapes as the figures and the concave ones as ground. This is independent of whether the convex shapes are black as in (a) and the concave ones white, or vice versa as in (b).

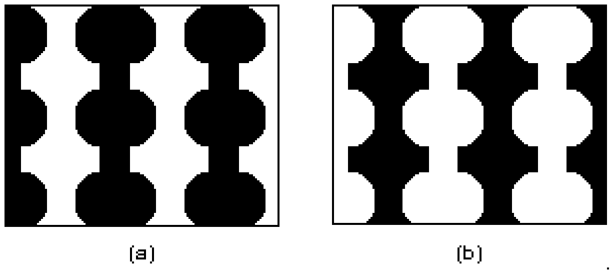

Fig 15 In general it is the convex shapes that are seen as figures against ground

To check whether the matter is influenced by degree of symmetry rather than convexity/concavity, Kanizsa and Gerbino (1976) went on to test subjects using diagrams similar to those in Figure 16 in which the concave shapes are symmetrical about the vertical and horizontal axes but the convex figures are only symmetrical about the horizontal axis. Once again the overwhelming majority of subjects saw the convex shapes as the figures whether they were drawn in black, as in Figure 16

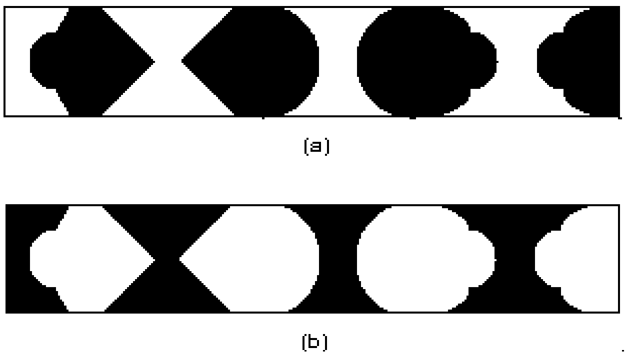

(a) or in white, as in Figure 16 (b).

Fig 16 Here again, the convex shapes are seen as figures on a ground of the opposite colour

This seems to confirm an innate preference for convexity in two-dimensional shapes. It is not clear, though, that this preference transfers to a preference for convexity in the third dimension (although this can probably be inferred).

As with the other grouping principles, the figure/ground, relative size principle can sometimes be overridden by a different preference: in this case for certain orientations (as in Figure 17 where, in both images it is the upright one that dominates).

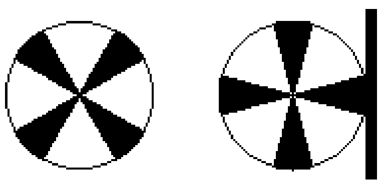

Fig 17 Our perception of what is figure and what is ground changes with the orientation of the figure suggesting that orthogonal relationships are preferred

Symmetry

We group together symmetrically arranged items and find it easier to make sense of symmetrical groupings than asymmetrical ones (Figure 18).

{ } { } { } { }

} } { } } { { }

Fig 18 Symmetry is a strong grouping principle. It is very much easier to make sense of the top line of the figure than the bottom one

It is known that we make less eye movements when dealing with symmetrical figures and it is probable that they take less cognitive resources to process.

Like all the principles, symmetry can also be upset by context. Thus the dot in the centre of the square in Figure 19 can appear not to be in the centre when an additional off-centre square is added.

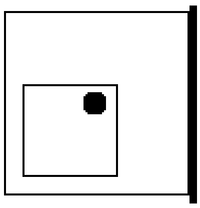

Fig 19 The addition of an off-centre square upsets our perception of the central dot

Familiarity and context

Familiarity with a scene and its context affects our grouping—sometimes bringing about substantial changes in our understanding of what we see (Figure 20).

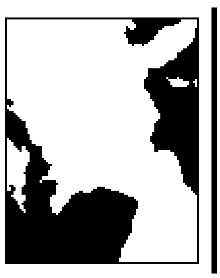

Fig 20 This familiar image is hard to recognise in this orientation. Turning the image through 90 degrees clockwise allows us to view it in 'correct' orientation

Our sense of 'objectness' often comes into play: knowing that we are looking at a partially occluded object immediately affects our ability to group. This is strikingly illustrated in Figures 21 and 22

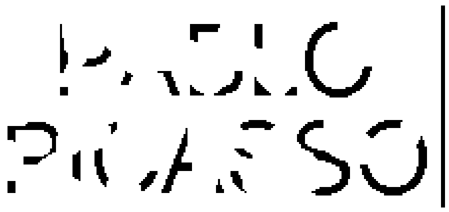

Fig 21 It is difficult to make sense of this image . . .

Fig 22 . . . until one realises it is partially occluded text

The context with which we frame a drawing is also significant. The squares and diamonds in Figure 23 are perceived differently according not, as might be supposed, to their relation to some absolute frame of reference, but to their placement within the drawing frame. This has obvious implications for the design of Virtual Environments, where the user's frame of reference—the sense for example of 'right way up'—is all too easily lost.

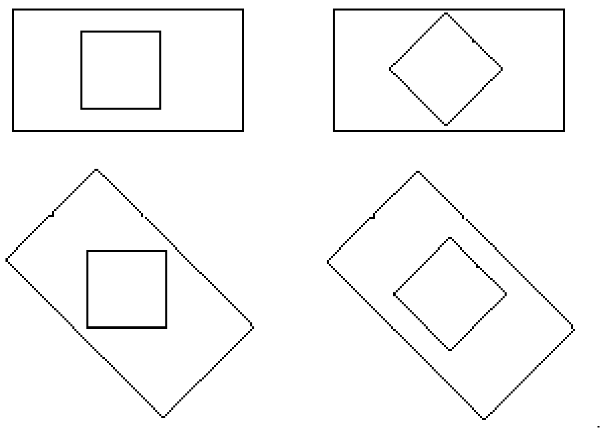

Fig 23 The relationship of the internal figure to the orientation of the frame determines whether
we will perceive it as a square or a diamond

Palmer (1992) has shown that context is also significant in allowing us to orientate ambiguously pointing shapes such as equilateral triangles. When seen alone, an equilateral triangle can appearto point in any of three directions (In Fig 24 a the triangle can be seen as pointing to 3 o'clock,7 o'clock or 11 o'clock). When seen grouped in company with similarly orientated triangles, all seem simultaneously to point in one or other of these directions (Fig 24b). The preferred direction of pointing changes when the triangles are aligned along an axis (Fig 24c). When they are aligned along one of their sides, they seem to point in a direction at right angles to the alignment (Fig 24d).

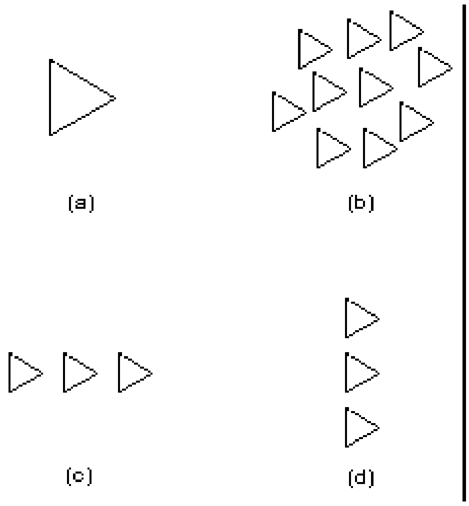

Fig 24 Equilateral triangles have ambiguous orientations that can be directed by appropriate grouping

Note the even more ambiguous effect of arranging similarly orientated triangles around a circle (Fig 25).

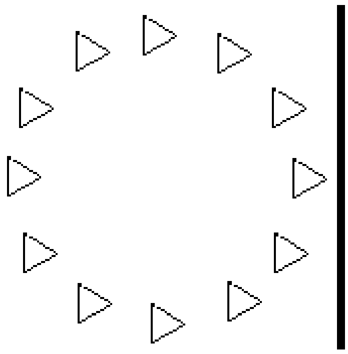

Fig 25 It is hard to perceive these equilateral triangles as having the same orientation

Common Fate

Items that move together are grouped together. This organising principle cannot be adequately illustrated in a still picture but requires animation or movies to be fully appreciated. One of the most impressive manifestations of the principle of common fate was performed by Johansson (1973a, 1975) who attached lights to joints of a black-clad actor and filmed him as he moved across a darkened room. When the actor was stationary, no pattern could be discerned in the lights. But, as soon as he walked, it was easily possible to identify the very sparse pattern as that of a moving figure. Pavlova (1992) has shown that, when children aged 3-5 years were tested with animated cartoons consisting of moving dots attached to the main joints of an invisible man and an invisible animal moving as if on a treadmill, the 3-year olds were able to recognise the moving displays and the 5-year old’s performance was as good as adults. Static versions of the display however were not recognised. Familiarity and context must also play a part here.

The active eye and the layout of information

It is clear that perception is a process which involves action on the part of the observer in relation to the observed. Even something as apparently trivial as decorative pattern engages the observer's mind in an active way. Gombrich (1979) argues that the faculties of expectation and extrapolation are essential to an appreciation of pattern, so that much of the pleasure of pleasing patterns lies in their touching the interplay between utter predictability on the one hand and chaotic unpredictability on the other. It is difficult to see how this could work unless, as Gombrich suggests, the mind is always predicting, and directing the eye to seek what it expects to find. The decorative artefact balances between redundancy, where no new information is introduced to the scanning eye and which the mind soon tires of attending to, and complete variety, which seems to the mind equally unrewarding of attention.

Gombrich is dealing with decorative surfaces, and uses the language of information theory to do so. However our need is often to deal with the conveying of information in the more limited everyday sense—the transfer of facts from author to observer. The structuring of information on a two-dimensional surface conforms to such obvious rules that we tend to take it for granted. Nevertheless such structuring can be regarded as providing useful meta-information which enhances the understanding of the content-information. So many design procedures make use of the spatial organisation of information in order to enhance understanding that only one or two examples need be cited.

Miles (1987) succinctly puts the case for spatial organisation in graphic design: 'Typography and layout, thoughtfully used, can provide signposts which reinforce the significance of each passage and help the readers to find their way about in a document.' A gamut of conventional techniques used in laying out a page relies on our sure interpretation of juxtaposition, separation, alignment and non-alignment, relative sizes and so forth. It includes indents, paragraph spacing, columns, tabular layouts, and of course 'white space'. The findings of the gestalt psychologists described above seem to suggest that this works at a basic precognitive level, though no doubt strongly reinforced by the learning of conventions.

Tufte devotes a whole chapter of Envisioning Information (1990) to another spatial artefact: layering and separation. Most of the book is devoted to the representation of multivariate information on a two-dimensional surface. In discussing layering, none of his examples depicts a third dimension in an obvious pictorial way. However, many of his examples allude to the missing dimension by means of tonal and colour differences. For example, in discussing the relation between signal and background—marks providing substantive data and those providing the framework in which it can be understood—he draws attention to the importance of suppressing the tone and saturation of the background. Perhaps our tendency to give attention to the marks of highest contrast as against those of lowest derives from our experience of the landscape, where tones and hues diminish with distance and we must attend for reasons of survival to those things which are nearest.

Interface design guidelines depend on the use of position on a surface for much of their effect. Examples (taken from Brown 1985) include:

- Use invariant fields on each screen: If an item of information has a similar function on similar screens, do not move it around. This advice denotes a use of space informed by restraint rather than virtuosity. Effort may be required by the designer to ensure that this consistency is achieved.
- Present lists in useful orders: Simply to lay out information in a linear array, so that items

are sequenced in a way which 'makes sense' in relation to their individual meanings, uses space importantly, if so simply that we often take it for granted.

- Arrange information in logical groups: Clustering (and aligning) objects so that those which 'belong together' are seen together enhances the meaning of the components. Such spatial organisation has the status of a grammar, like that which gives words their significance in a sentence. An example familiar to all computer uses is the GUI convention of clustering 'radio-buttons': those within a set are mutually exclusive while there is no exclusivity between sets. The user can only make use of this knowledge if it is clear which set each button belongs to.
- Minimise pointer movement: Many early hypermedia products imitated the conventions of turning the pages of a book, by placing the 'next page' and 'previous page' triggers respectively at the right and left of the screen. This was a blunder significant for any discussion of Virtual Reality, since it was based on an ill-considered transfer of some attributes of real paper documents to a virtual environment. It involved users in wasteful pointer movement whenever they wanted to 'flip pages' back and forth. Soon afterwards designs began to appear which still borrowed the idea of the right button denoting 'next' and the left button 'previous' but which brought them together in a tight cluster of controls. Taking liberties with reality had proved the better solution than following it slavishly.

The significance of these guidelines is that they are about constructing meaning—or at the least, amplifying meaning—using space. Of course, information structures of this kind can be enhanced by the availability of an additional, third, dimension. Even so, it is worth remembering that when using three dimensions in this way, three will often not be 'enough', just as two was not (Information spaces: from database to the Virtual Museum page 69).

Movement perception

The visual perception of movement is important to virtually every living species. Indeed even organisms without vision usually have sensors to detect movement. Obviously detection of movement plays a vital role in the survival of animals: they must be good at perceiving movement of predators and of likely prey. Inability to do this would result in disaster and it is often more important to detect immediately that something has moved rather than to know straight away what that something is (or even in which precise direction it has moved). Sekuler (1975) proposes that:

> During evolution, motion perception was probably shaped by selective pressures that were stronger and more direct than those shaping other aspects of vision. . . As a result of such selective pressures, our visual systems contain neural mechanisms specialised for the analysis of motion (p385).

The importance of movement perception

We have at least two needs for motion perception: one, to make sense of the world as we move through it—self-motion or 'egomotion'—and, two, to understand objects that move about us. Much of the time, of course, both these needs come into play simultaneously. However, it is puzzling how easily we can normally distinguish between the movement of objects and our own movements although, as far as the retinal image is concerned, there does not seem enough information to do this. Note, too, that the retinal image is also altered by movements of the eyes or head. Harris, Freeman and Williams (1992) postulate perceptual mechanisms that would, if they exist, help to explain how we are able to separate out these different retinal effects. Cutting et al (1977) suggest that movement detected in parts of the retina at the periphery of vision allows us to perceive self-motion whereas the same stimulus at the centre of vision allows us perceive object motion.

In their comprehensive review of aspects of visual motion analysis from a computational perspective, Hildreth and Koch (1987) point out that:

> The pattern of movement in a changing image is not given to the visual system directly, but must be inferred from the changing intensities that reach the eye. The 3-D shape of object surfaces, the locations of object boundaries, and the movement of the observer relative to the scene can in turn be inferred from the pattern of image motion. Typically, the overall analysis of motion is divided into two stages: first, the measurement of movement in the changing 2-D image, and second, the use of motion measurements, for example to recover the 3-D layout of the environment. It is not clear whether motion analysis in biological systems is necessarily performed in two distinct stages, but this division has served to facilitate theoretical studies of motion analysis and to focus empirical questions for perceptual and physical studies (p480).

The need for good movement perception

Nakayama (1985) suggests seven possible reasons why movement perception is important. These are:

to enable us to

- derive the third dimension
- calculate time to collision
- distinguish figure from ground
- ascertain information about our own movement
- stimulate eye movements
- understand pattern
- perceive moving objects

Warren (1995) goes further and stresses the importance of movement and action to perception as a whole:

> Traditionally, the problems of perception and action have been treated as logically independent. It has been assumed that the goal of perception is to recover objective quantities such as size, distance, shape, colour and motion, yielding a general-purpose description of the scene that can provide the basis for any subsequent behaviour. Considering vision in the context of action has important implications for this view of perception (p264).

He goes on to show that some of the anomalies of perception that are apparent when we consider it in a still 'snapshot' way unrelated to action may not matter so much when we see perception as something that unfolds over time. Viewed thus:

> . . . judgments at any instant may be qualitative or even nonveridical, and yet over the course of the act adaptive behaviour emerges from the animal-environment interaction (p264).

All this points to the fact that the detection of motion is highly significant to us. It is all the more surprising therefore that research into this aspect of our visual perception mechanisms seems to lag behind work on other aspects. While colour vision and stereopsis have received the greatest attention by researchers, as Nakayama (1985) says,

> . . . it is clear that colour processing is not present in all species and that binocular vision is restricted in animals with laterally placed eyes. As such numerous animals either lack colour vision or significant binocular vision or both. No animals have been found that lack mechanisms for motion processing (p627).

Development of motion perception

Referring to the growth of space and movement perception in children, Kellman (1995) tells us,

> From the earliest ages, motion attracts attention, and infants orient toward moving stimuli by using head and eye movements . . . The causal direction of the connection between motion and information is not known. Infants might be hard-wired to attend to moving things—a useful adaptation for learning about objects and events. Alternatively, it may be information, not motion, that guides attention. Infants may preferentially attend to events more than static scenes because more or better information about spaces and objects is available to them from kinematic sources than from static ones (p346).

It appears from studies by Kaufmann (1995) that a child's ability to detect very rapid motion is fully developed soon after birth whereas the ability to detect very slow motion seems to improve gradually with age. As adults we can detect 3-6 mm of movement per second of an object when it is 1 m away. However, it appears that we substantially underestimate the velocity of moving objects (Caelli 1981 pp145-171).

Ullman (1987) suggests that the importance of motion detection has led to particular physical results:

> In view of the central role of motion perception, it is not surprising that the analysis of visual motion is wired into the system from the earliest processing stages. In some species, including the pigeon and the rabbit, rudimentary motion analysis is performed as early as the retinal level (p1280).

We know too that, after recovery from damage to the human visual processing mechanisms, it is our perception of movement that returns first. Furthermore, unlike for some other aspects of vision, most of the retina is available for motion detection and the thresholds for discovery are more or less the same over the whole area (Krumhansl, 1984).

The mechanisms of movement detection

In view of the significance of motion perception and the different forms it has to take, it has been postulated that several mechanisms are at work. For example, Juola and Breitmeyer (1989) tell us that:

> The sheer variety of motion phenomenon that can be perceived argues for the necessity of having more than one detection system, perhaps located at different levels in the nervous system. It is not the case that an image moving across the retina is sufficient stimulus for the perception of movement. For example, the retinal image moves whenever a saccadic eye movement is executed, yet we are unaware of any motion. Conversely, we can track a moving object against a featureless without necessitating changes in the retinal image, yet a strong sense of movement results (p251).

As with other aspects of visual perception it is possible to locate specialised areas. Anderson (1989) points out that:

> A substantial body of evidence suggests that visual motion analysis is treated by specific brain regions. There are several accounts of brain lesions in humans that produced deficits in motion perception without deficits in other forms of vision (p383).

Anderson locates the 'pinnacle of a prescribed hierarchy in motion processing' in the posterior parietal cortex—Brodman areas 5 and 7. Zeki (1990) confirms that, in macaque monkeys:

> The 'funnel' for the motion pathways of the visual cortex is area V5 . . . All its cells are responsive to motion in the field of view and the overwhelming majority are directionally selective. . . The fact that none is wavelength selective and the majority are not orientation selective either . . . led me to propose that it is a cortical visual area specialised for the detection of motion in the field of view (pp321-322).

From their examination of healthy and brain-damaged patients, Schenk and Zihl (1995) conclude that there are two distinct mechanisms involved in motion detection. One of these is concerned with the analysis of global motion. The other deals with extracting form from motion.

In their discussion of human motion perception and the physiological nature of motion detectors, Burr and Ross (1986) point out that accurate perception of form in motion involves visual integration of a type that eliminates smear. Image motion does not produce the same problems for the eye as it does for the camera. People can clearly see objects in motion, and they can see motion on cinema or TV when what they are shown is actually a sequence of stills. They look at mechanisms that allow for understanding when there is real motion, and interpolation when motion is sampled as in cinema. They suggest that such visual abilities may be explained by specialisations of visual neurons.

In examining moving images of a Mondrian-like nature, Zeki and Lamb (1994) conclude three things:

- that an image of the visual world is not impressed upon the retina, but assembled together in the visual cortex;
- that separate attributes of the visual scene are processed in geographically separate parts of the visual cortex, before being combined to give a unified and coherent picture of the visual world;
- that the attributes that are separated, and separately processed, in the cerebral cortex are those that have primacy in vision. These are colour, form, motion, and possibly depth.

Using PET scanning DuPont et al (1994) confirm that several brain areas are involved in the perception of movement. In addition to sites on either side of the brain at the border between Brodmann areas 19 and 37, a V1/V2 focus and a focus in the cuneus, they observed activity in other visual areas in the cerebellum and in two presumed vestibular areas—at the posterior part of lateral sulcus and at the border of Brodmann Areas 2 and 40 (Figure 1).

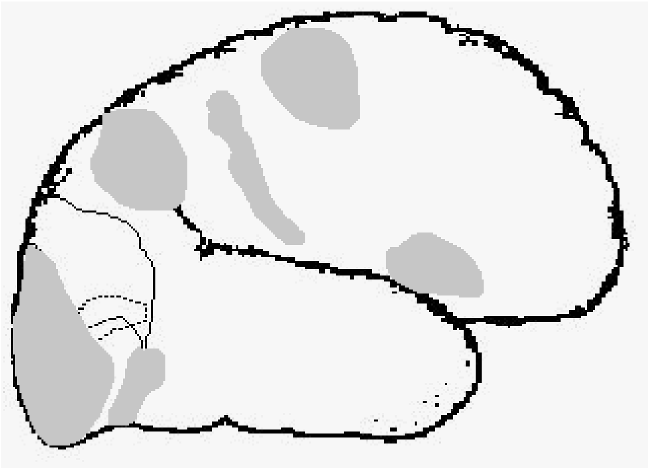

Fig 1 Diagrammatic view of the brain areas identified by DuPont et al (1994) as being involved in the perception of movement

Early work

In their overview of motion detection research, Smith and Snowden (1994) suggest that the work of the Austrian scientist and colleague of Helmholtz, Siegmund Exner (1846-1926), provided the starting point in the latter part of the nineteenth century.

> He demonstrated that motion could be perceived from two stationary images (sparks of electricity in his case) presented in quick succession—a fact exploited by television and 'movies'. This had been known for some time previously . . . However, Exner's great insight was that this perception of movement could be elicited from two sparks that were so close together in space that they could not be distinguished. Under these conditions it seems impossible that the observer could infer (consciously or unconsciously) motion from a knowledge of position and time. It therefore follows that motion perception must be a sensation in its own right, not one derived from a sense of position and time (p5).

Previous to Exner's studies, others had commented on aspects of motion such as optical illusions but little detailed study had been carried out. However, the stroboscope—for examining objects in motion—had been in existence since the 1830s (see Boring, 1942 pp588-596). About that time, too, Robert Addams (1834) reported on a perceptual anomaly about motion that he had encountered on a trip to Scotland:

> During a recent tour through the Highlands of Scotland, I visited the celebrated Falls of Foyers on the border of Loch Ness, and there noticed the following phenomenon. Having steadfastly looked for a few seconds at a particular part of the cascade, . . . , and then suddenly directed my eyes to the left, to observe the vertical face of the sombre age-worn rocks immediately contiguous to the waterfall, I saw the rocky surface as if in motion upwards, and with an apparent velocity equal to that of the descending water, which the moment before had prepared my eyes to behold this singular deception.

In parenthesis it is worth noting the style of language used in this description of what has now become known as the 'waterfall illusion'. It was written for a scientific journal and yet the style does not differ markedly from the sort that was used a few short years later by, for example, John Ruskin (1819-1900), in his descriptions of natural phenomena in his books on art. The differences we now see between scientific reporting and writing on art seem to have come about in this century: the two cultures of CP Snow are clearly a modern concept.

Despite some earlier work and remarks on motion phenomena, it was not until the late nineteenth century that systematic studies of motion occurred. We should also remember that it was in the period 1870-1890 that the English photographer, Eadweard Muybridge (1830-1904) in the US and Etienne-Jules Maray (1830-1904) in France, carried out and published studies on the movement of humans and animals. As well as leading in the mid-1890s to the development of the motion picture, these studies must have also spurred an interest in the phenomenon of movement perception.

The optic array and optic flow

Many believe, with Gibson, that the major element in motion perception (and, perhaps, in visual perception generally) is the so-called 'optic flow' (Figure 2).

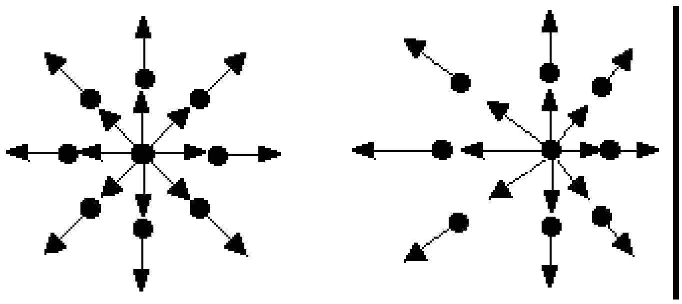

Fig 2 On the left is the diagrammatic situation when we approach a surface at right angles to us. On the right, when we approach a surface at an angle, with the left edge nearest to us

As Harris (1994) perceptively comments:

> As we move about and scan the visual world, the images of things move about, changing their relationships in a complex dance. JJ Gibson . . . was the first to understand this dance . . . Probably Gibson's greatest contribution was to redefine the dance-floor, emphasising the amount of information potentially available to an observer in the transforming optic array rather than the instantaneous fragments provided by a pair of retinal images.
>
> The optic array is the three-dimensional bundle of light rays that impinge from all directions upon each point in an illuminated world. Objects in the world can be thought of as labelling specific rays, so producing a global pattern of light intensities. A retinal image provides access to only part of the optic array at any one time, but a stationary observer can sample different parts by eye movements and head rotations. By changing position, the observer can sample the different optic arrays impinging on neighbouring points in space. However, sampling in this case should not be thought of as a discrete process. Rather, as the observer gradually moves, so each ray gradually moves, thus producing the smooth transformation in the optic array that Gibson called the optic flow (p307-308).

Gibson believed that we 'picked up' on the optic array and optic flow in order to understand the world. Unfortunately, he proposed no explanation as to how the 'picking up' might be achieved. Johansson (1994) agrees with Gibson and suggests that the principles for decoding the optic flow are 'built-in to the specific visual systems of the species' (p311). He outlines his approach to decoding the optic flow and what he calls the 'optic sphere' in Johansson and BÜrjesson (1989).

Warren and Hannon (1988) conclude from studies of computer screen images that optic flow is sufficient at least for perceiving the direction of self-motion.

Many workers have made proposals for how we might cognitively deal with optic flow. Harris, Freeman and Williams (1992), already referred to, develop a concept by Koenderink and van Doorn (1976) who treat optic flow as a vector field in which the operations of translation, div (divergence), curl (rotation) and def (deformation, shear) apply (Figure 3).

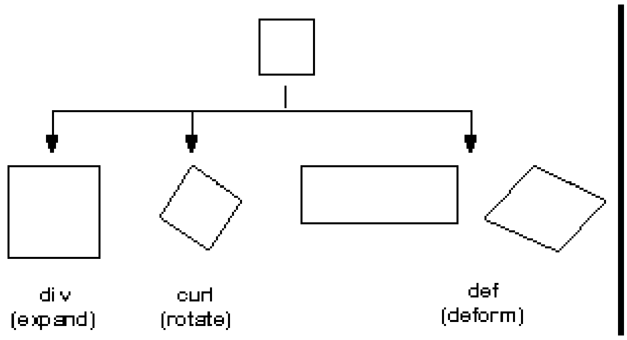

Fig 3 showing the vector field operations of div, curl and def

De Bruyn and Orban (1990) give evidence to suggest that we are able to detect div, curl and def over much of the field of view.

Div, curl and def translate into operations on the optic flow lines as in Figure 4.

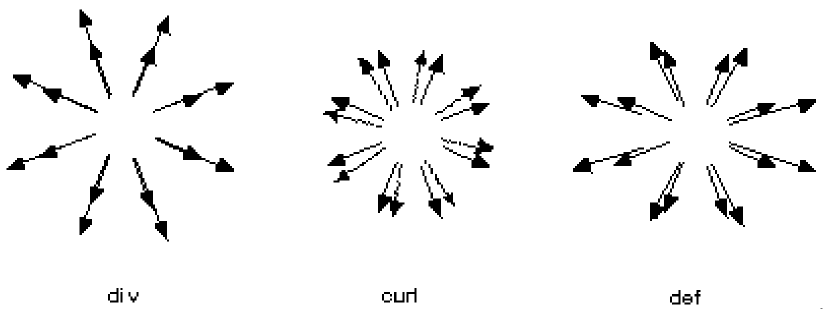

Fig 4 showing how div, curl and def affect the optic vector field

Dodwell (1983), basing his work on that of Hoffman (1996), uses the vector field ideas of Lie algebra to show how, with appropriate perceptual mechanisms, we might simply process optic flow (see also Hoffman 1970; Hoffman and Dodwell 1985). However, Caelli (1981 p143) takes the view that Hoffman's formulation is a meta-language that allows us conveniently to talk about some aspects of perception rather than giving us an explanation of it. He suggests that lower level models are needed before we can understand the situation. Harris, Freeman and Williams (1992) try to provide such models.

Biological motion detection

A curious ability we possess was first demonstrated by Johansson (1973a) who showed that, if small lights are placed on the ankles, knees, hips, wrist and elbows of a black-clad actor, and a movie is made of the person walking in a dark room where only the lights are visible, we can easily understand what we are seeing. Conversely, we cannot understand what we are seeing if no movement takes place.

This phenomenon has been studied by a number of workers. Kozlowski and Cutting (1977) show that the gender of an actor can be deduced from the moving pattern of lights. Runeson and Frykholm (1983) show that people are remarkably accurate in guessing the distance that the black-clad actors are able to throw small sandbags even though unable to see the bags or anything more than the lightspots on the actors' joints. Pavlova (1992) showed that quite young children are able to recognise the patterns—Berthenthal et al (1985) estimate that the ability to judge biological motion develops at around 6 to 9 months. Mather and West (1993) show that animals can also be recognised from this apparently limited information source. Sumi (1984) shows that, even if the film is run backwards and upside-down, subjects still recognise a human walker. What is more they perceive it as a walker with a peculiar gait rather than as an inverted image of a person moving backwards. Sumi concludes that this conception seemed to arise from the fact that the actor's arms were perceived as legs and vice versa. Dittrich (1993) shows that recognition of what is happening is more readily achieved when the figures are waking or running than carrying out social or instrumental actions. Further, that recognition is not greatly impaired if the lights are placed not on the joints but between them.

It is probably our ability to detect biological motion that explains the results of studies by Campbell (1979) who found that, in getting young children to understand still drawings of running and walking figures in implied motion, it was more important to show the figures in active postures than to employ motion cues like speed lines such as cartoonists sometimes use. Not surprisingly, Friedman and Stevenson (1975) found that older children were able to understand both postural cues and the cartoonists' conventions. They too, however, seemed to favour active postures.

Movement in graphics and virtual reality

There is little doubt that movement helps in our understanding of objects. Frequently, drawings that make little sense when shown as stills are understandable when they are animated. This is particularly the case in viewing 'wire-frame' drawings. Wire frame drawings of three-dimensional scenes often contain too many lines for us to disambiguate when seen in still form. When animated to allow for panning across the scene and rotating objects, however, they spring to life and the third dimension appears effortlessly. Often, indeed, a wire-frame animation seems to convey more information about the form of an object than does a movie of the same thing (although, of course, it cannot tell us anything about the surface properties of the object).

The attractions of three-dimensional drawings, particularly animated three-dimensional drawings, are such that, sometimes, they are used inappropriately—for example, to illustrate two-dimensional data. When this happens not only do they not enhance the data set, they sometimes falsify it. We should also note that, although we seem to have separate mechanisms for dealing with motion, colour, and form (at least), it is clear that we can better understand subtle changes in a scene when we attend to one of these aspects only (Corbetta et al 1990). From this, and other studies, we would conclude that, if we wish to have someone focus their attention on a particular element in an image or animated scene, then we should not create conflict by changing motion, colour and form together. This would be especially the case if abstract data is being shown.

Interestingly, in an admittedly slightly limited and uncontrolled study by Felix (1995), we find that there is no evidence that being, as it were, inside a virtual space and walking about it, as opposed to viewing it from outside from different viewpoints, necessarily improves understanding of it. In both cases, motion assists. But immersion in the space does not necessarily provide good understanding of it.

Summary

The detection of motion is a strong element in visual perception and we gain much information from it. Indeed, some believe motion to be the prime source of our understanding of the world. Sometimes, it is only through motion that we can disambiguate conflicting signals and it is likely that our ability to process optic flow is primal.

Mnemonic uses of space

There is a historical use of space previously known only to a small group of specialist scholars but more recently influencing a number of designers of information media. This is the Art of Memory, a spatial mnemonic system. Yates' classic book on the subject (1966) provides a fascinating in-depth account, but we can summarise the main points of interest here.

The assumption of the Art of Memory is that we are predisposed to remember things in the context of place, even where there is no significant connection between the thing remembered and the place where it is located, so that recalling the space is a powerful trigger to the recall of the associated information.

According to Cicero's De Oratore, the poet Simonides invented the Art when called upon to name the unrecognisable victims of a physical disaster—the demolition of a building full of dignitaries from which Simonides himself escaped through the intervention of the gods. He was able to name the victims by recalling where they had been seated.

Such spatial mnemonic techniques were used subsequently during the Middle Ages (it has been argued by Yates and others that the Cathedrals were organised as aids to remembering the scriptures) and later, in the Renaissance, by such as Giulio Camillo (1480-1544), Ramon Lull (1235-1316, his works revived in the 15th Century), Giordano Bruno (1548-1592?), Peter Ramus (1515-1572) and Robert Fludd.

Many practitioners of the Art combined what we would now regard as magical nonsense with what we would now regard as fruitful cognitive theories. As with the present-day memory performer on stage or television, the techniques were guarded secretively in order to enhance their impressiveness. This obfuscation tended to lead to the denigration of the theories with the rise of scientific rationalism, which proposed an alternative model (see below), in the 17th Century.

The basic method of the Art is to imagine a space, perhaps schematic but usually architectural, which contains the various things to be remembered. Specific ideas can be contained within other more general ideas, by virtue of being in niches within rooms. Rooms may be badged by the statue of a saint, for example, representing some important principle. Some advocates of the Art devoted much energy to devising structures which in themselves had special meanings, so that the shapes of room and the ways in which they are connected take on semantic importance—a three-dimensional semantic net, perhaps, but made memorable by being given form as an imaginable physical building.

Importantly for any consideration of VR is the fact that, rather than simply imagining these architectural structures, most of the leading mnemonists had plans to actually build them! Camillo built a memory 'theatre' described thus:

> They say that this man has constructed a certain amphitheatre, a work of wonderful skill, into which whoever is admitted as a spectator will be able to discourse on any subject no less fluently than Cicero.
>
> Viglius Zuichemus, writing to Erasmus, 1532, quoted Yates 1966, p135

The Art seems to make use of two or three aspects of spatiality:

1 Recalling a view of a space is easier than recalling abstract symbols (such as abstract concepts or pieces of language). Concreteness produces memorability.

2 Spatial constructions, for example a building or the human body, have a coherence and logic to them which can be used mnemonically to connect one idea to another. Interestingly, the Art begins to decay at the time that taxonomies become a prominent tool for thought (Foucault 1974 p125 passim). In a taxonomy, for example a zoological family tree, the structure arises out of the subject matter, rather than as in the Art of Memory being arbitrary.

One possible advantage of architectural models of conceptual structures is that the omission of any part leaves an obvious 'gap' in the structure. For something organically random (like a semantic net) there is little help for the user who forgets some major chunk of the structure (a uniformly branching coral is not structurally odd if part of it is missing), whereas when information is mapped to a finite, structured environment, omissions are obvious.

Against this we must set the disadvantage that our information structures are now routinely too complex to be mapped onto memorable architectural spaces.

3 We can imagine ourselves taking a route through a structure, as a way of converting a 3D environment into a linear sequence. In principle, it should be possible to take different 'tours' of the same space of ideas. The proposers of memory theatres, whether real or imagined, thought of every possible trick to embed their ideas in the user's memory, including the use of colour to distinguish one part from another, the use of distinctive columns rather than a uniform design so that one area will not be confused with another, the use of magic numbers—presumably their emotional charge made them more memorable, even though the reasons they were used were directly occult!

In practical terms, what has the Art of Memory to offer us in the design of virtual environments? A number of the debates which polarised academic communities in the Renaissance are still of relevance today:

Imaginary or real structures
Should the user imagine a building which was purely ideal, or use the remembrance of an actual building? We could represent ideas and the connections between them using abstract models, or we could project information structures onto reconstructions of real structures. Advantages of a familiar structure would seem to include not only the greater likelihood of users remembering the information when they leave, but also the corollary advantage of them knowing how to find their way about the place when they are 'in' the Virtual Environment. Of course we do not need a universal principle: each Virtual Environment can be built according to need.

Memorable imagery
'Images must be lively, active, striking, charged with emotional affects so that they may pass through the door of the storehouse of memory,' says Giordano Bruno in 1591 (Yates 1966 p286). Nothing in the intervening 400 years leads us to believe that Bruno was wrong, and if we are to construct Virtual Environments with educational purposes in mind, we need to ask ourselves what would constitute the lively, active, striking and emotionally charged equivalents for our own time.

Arbitrary or semantic structures In 1584 a row broke out between the followers of Ramus and of Bruno. With hindsight, we can see that the Ramists would eventually triumph. They advocated a 'dialectical' system in which each broad category of knowledge was repeatedly subdivided—the form of spatial organisation familiar to us as a binary tree. Ramus held that this structure absorbed the art of memory into that of logic. It is an ancestor of the taxonomic approach which we now in many ways take for granted, the basis for example of the Dewey system for organising knowledge in libraries. As a Protestant, Ramus had a strong antipathy for representational images, but still valued the spatial laying out of schema as a graphical framework for his dialectical system. He particularly abjured the use of imagination, and the assistance to memory provided by striking and stimulating images which had been a mainstay of his opponents' Art.

Special shapes for greater effect; symbolic architectures

Much of the literature of the Art concentrates on the literal shapes of knowledge, in some cases architectures buildable within the technology of their time, in other cases ideal structures involving for example intermeshing circles, or even the human body. Even after the obsession of some adherents with magic numbers, planetary harmony etc is stripped away, it is clear that these forerunners of ours had given a very great deal of thought to the shape(s) of knowledge. They would have been intrigued, probably amused, to see the 'virtual museum' of the Dorling Kindersley eye-witness series of encyclopaedic CD-ROMs. In some cases, their own systems (for example Bruno's) were intended to house all knowledge.

Memory—perhaps enhanced by memorable spaces based on some of the ideas of the Art—has a great deal to offer in terms of assisting learners to find information. For example, experiments have been undertaken using 3D spatial organisation to assist users in forming cognitive maps of information (Shum 1990). By improving their abilities of recall, users may be helped both to recall information when not using the system, and in refinding information in the system on future occasions. The whole subject of wayfinding in real environments—and the role of memory in the process—has been discussed by Passini (1992).

It is interesting to note that memory has become almost an occult skill in our own education system. For example, the idea of rote learning is deeply unfashionable, but perhaps the deprecation of discredited methods for instilling memory skills has accidentally spilled over into a deprecation of memory as such. Books on memory are regarded with embarrassment by most academics and are felt not to be part of serious scholarship. Such an attitude would have astonished most previous centuries, and maybe we should begin to question it.

Social interaction mediated by space

The spaces of the Art of Memory were designed in relation to the single observer, in many cases only existing inside the head of the individual. The issue of spaces as a place of mediation between one individual and another did not arise. Indeed at that time, architecture itself was not seen as the creation of spaces as such, but rather of structures. The idea of studying space and the way in which it is used socially had to wait until the twentieth century (see also Architecture as spatial design, page 39).

Influenced by the work of Whorf, Edmund Hall (1959, 1966) drew attention to the uses of space in mediating social interaction at both a biological and cultural level. His guiding principle can be summed up in his own words:

> Culture hides much more than it reveals, and strangely enough what it hides, it hides most effectively from its own participants.
>
> Hall 1959 p30

The idea that different cultures interpret space and location in different ways is now widely accepted, even if the related notion that we actually perceive in different ways is by no means granted (discussed above, page 7). The design of virtual environments would ignore this at its peril.

In his attempt to reveal how we communicate by our behaviour, Hall identifies a number of ‘primary message systems’ which define the way a society works: interaction; association; subsistence, bisexuality; territoriality; temporality; learning, play, defence and exploitation (use of materials) (Hall 1959). Spatiality is fundamental to many of these, and all have a material aspect of some kind. These were the themes of Hall’s second book, The Hidden Dimension (1966) which is still worth reading. Hall looks at distance regulation in animals (including humans), observing the constants across different societies and the culturally determined variations. He studies crowding and social behaviour in the light of what he defines as intimate distance, personal distance, social distance and public distance. Very unusually, he looks (though not in depth) at auditory space and olfactory space.

The single most important point about Hall’s work is his emphasis on the silence of the ‘silent language’, and the hiddenness of the ‘hidden dimension’: in other words, our inability to readily see the biases and limitations of our own culture. This argues for investigating widely the possibilities which other cultures (or other views of our own) might offer, in our case in designing environments and objects for social interaction.

Spaces and power

It is not only in the work of anthropologists that the social uses of space are described. Novelists and other writers have made extensive use of spatial meaning, while postmodern geography has interesting things to say about the many shapes of space.

A gamut of different texts and images has made use of spatial organisation as a social and political metaphor. One of the most obvious is the mapping of social strata to physical levels. From Verne, through Wells to Fritz Lang’s Metropolis and Blade-Runner there has been a recurrent use of the idea that a social layering of society can be represented by a physical layering of its people, usually of course with the ‘lower orders’ beneath. However, some fictions have placed the ‘upper classes’ in underground utopias, protected from climatic disaster, environmental pollution and so forth, while the ‘lower orders’ are left above, outside on the exposed surface. Many of these are studied with wit and care by Williams, in Notes on the Underground (1990).

Williams’ work, like so much in this report, emphasises how spatial structures, be they built or excavated, real or fictional, cannot be considered independent of the messages which they convey. Giedion, quoted by Williams, claims that from Roman times onwards there was an urge to construct awe-inspiring space of great dimensions: the concept of a hollowed-out interior. This passion for sublime immeasurability unites spaces as diverse as the Pantheon, Piranesi’s ‘prisons’, Coleridge’s palace of Kubla Khan...

> Where Alph, the sacred river, ran
> Through caverns measureless to man
> Down to a sunless sea...
>
> It was a miracle of rare device,
> A sunny pleasure dome with caves of ice!

... to the deep space of Star Wars and many computer games.

Williams documents a largely ignored phenomenon, the fear and loathing in which underground space was held until the 19th Century, when such emotions were replaced by a fascination with the idea of subterranean life and society. ‘Once underground however [in these fictional subterranean communities], society proves vulnerable to other catastrophes—not natural disasters, in most cases, but social ones arising from humanity's inability to live harmoniously in an enclosed environment.' So according to Williams, an important concomitant of the idea of social layers was the idea of transgression, when one group invades the space belonging to another. We can look at these 'politics of space' in the familiar context of buildings.

The meaning of buildings

> There is no a-spatial society and no a-social space... Society is organised in a way which can be described in the abstract but which, in the material world, is embedded in space.
>
> Markus 1993 p13

Markus studies the effects that principles of social organisation have had on the design of actual buildings. In Buildings and Power (1993) are catalogued the asylum, prison, workhouse, mill and school, among others. Highly influential was the Panopticon: a building designed to provide an overseer with maximum views of all the activities under his control, for which its designer Jeremy Bentham claimed 157 advantages, including 'morals reformed', 'health preserved' and 'industry invigorated'.

The most important issues that Markus' work points to—which we must consider in relation to Virtual Reality—is that of control: what spaces do to people. As we cease to see the experience of VR as a solitary affair between the individual user and the space, and increasingly as a shared environment supported by networked computing, this becomes even more important. Most of the buildings which Markus documents have the form they do because in them someone is controlling someone else, even through the simple act of vision. In some buildings, many things can be seen simultaneously by all visitors, eg the Crystal Palace, while in others viewing is strictly limited, and perhaps unidirectional, as in the harem. One building is designed so that the inmates may see the instructor or preacher but not each other; another is designed so that all present may be viewed from a control position, but not themselves view. In case this seems a purely historical phenomenon, it is worth mentioning a recent project to create a pan-European electronic publishing project which, under the control of a newspaper magnate, was designed to give him a view of the various distributed editorial workstations around Europe, but not provide the editors themselves with a similar view (DIMPE1990). Who can see what and under what circumstances is a vital issue. Interest in this aspect of space seems only to have emerged in the nineteenth century, perhaps because it was only then that the concept of privacy, the unacceptability of people looking into other people's spaces, really emerged. For example, the horror in Henry James' The Turn of the Screw is almost entirely based on inappropriate looking, with constant reminders that social and sexual propriety can be undermined simply by the act of looking itself. New permissions and prohibitions on seeing can be expected to emerge in Virtual Reality: it is too soon to say what form they may take.

What traversals are possible from one space to the adjoining spaces? For a simple matrix of nine rooms, many different topologies are possible. Markus adopts the Space Syntax methods of Hillier and Hanson (1984), which in themselves may prove useful in the evaluation and planning of virtual realities. These allow Markus to, for example, analyse a health centre to reveal that while the doctor needs to pass though only a small number of rooms to get from the staff entrance to his/her surgery, patients must pass through a minimum of six rooms go get to the surgery from the public entrance (Markus 1993 p13). Thus are the 'politics' of the organisation embodied in its architecture: it is not for nothing that the word 'accessible' has a metaphorical meaning as well as a literal one.

Of course even if spaces are physically connected, there may be prohibitions on who may use which routes. Girouard documents how in the Victorian country house 'an intricate system of backstairs and back corridors ensured that housemaids could get up to the bedrooms, dinner to the dining room and the butler or footman to the front door with the least possible chance of meeting the [owner's] family on the way' (Girouard 1978 p285). At an earlier period it was normal for a personal servant to be housed in the same space as the close-stool: 'the servant, the contents of the close-stool, and anything that was undesirable or private could move or be moved up and down the backstairs, preferably to offices in the basement' (ibid p138). What behaviours do these prohibitions and permissions promote? In a virtual reality any traversal may be forbidden to any user if we wish, but equally, the barrier created by something as apparently impenetrable as a wall may be permeable in a non-real world, if we wish it to be, as Carroll did for Alice when he allowed her to climb through into Looking Glass World.

What separations of people, and what confluences, do environments create? And what about the discreetness of the various virtual environments which are created, without thought for their interconnection? So far, VRML has hardly countenanced the issues of connecting one space to another, so that for example the initial opening of a site will always by default open it at the same entrance-point.

Some trajectories through buildings will have heavy traffic which moves rapidly, and otherwise will be almost empty, while others will be places specifically of relatively static congregation. A student project in the Centre for Electronic Arts was a commission to design a point-of-information system for a major museum, but it fairly soon became apparent that the client actually wanted a system that would tempt people to explore the furthest, neglected, parts of the museum. The last thing that was needed was a system which would encourage users to linger where the information machines were. Strategies for enticing users to explore neglected parts of a space have an obvious analogy in learning and learner-centred research.

The concept of affordances has been borrowed from Gibson's work on perception by design theorists, for example Krippendorf (1989). It refers to the way in which any designed artefact has (or should have) preferred 'readings'. So, for example when we look at a simple tool like a spade we are almost involuntarily drawn to the idea that one part, say the handle, is for our hands to hold, that another, the blade, is for piercing and cutting, and so forth. If we design well, what we have made should 'speak' its purposes and modes of operation. Architectural structures also have affordances. So for example, it is difficult to consider a Gothic cathedral without giving primacy to the experience of moving from the West door along the nave to the altar at the East, despite the fact that in theory the building could be appraised from any direction we wish. The application of the concept of affordances to the design of Virtual Environments is followed up below, page 52.

Architecture as spatial design

Spatiality has its own history. There is reason to believe that space has not always meant the same, even to those who made creative use of it. In his An Outline of European Architecture Nikolaus Pevsner (1943) tells us:

> What distinguishes architecture from painting and sculpture is its spatial quality. In this, and only in this, no other artist can emulate the architect. Thus the history of architecture is primarily the history of man shaping space...'

But it was not until the 1880s-1890s that 'architecture as space-making' became a theme for definitive discussion. Prior to that time architectural theory concerned itself more with the elements that shape enclosure rather than explicitly with the space created. Architecture was often seen as a sort of sculpture that you could enter or as 'frozen music'. Of course, we are referring here to the explicit notion of architectural space. Clearly, architects had always concerned themselves implicitly with spatial matters. Rooms were proportioned according to spatial considerations. Rooms of contrasting plan area and volume were intermixed so that one experienced a changing spatial impression when moving from room to room. And the way in which various elements 'modulate' architectural space was always taken implicitly into account. But the notion of an explicit 'architectural space' is largely a twentieth century notion.

Architectural space is modulated and defined by the elements surrounding and within it. Gottfried Semper (1803-1879), a German architectural historian, possibly the first theoretician to address the question of architecture as space directly, says 'The wall, partition or screen is that architectural element that formally represents and makes visible the enclosed space as such, absolutely, as it were, without reference to secondary concepts' (Semper 1989).

These modulating elements need be neither complete nor opaque. Thus a pierced screen or a row of columns will delineate a space even though these items will do no more than interrupt visual continuity. It should be noted that, the degree of visual opacity provided by a row of columns or a pierced screen varies in accordance with the sharpness of the angle at which it is viewed. The effect of this is that as one moves about a space surrounded by columns, one gets differing impressions of opacity and transparency, of enclosure and openness.

Space is also delineated by lowering a ceiling over an area or by changing its floor level with respect to its surroundings—although, significantly, lowering the ceiling has a seemingly stronger effect than lowering the floor.

When an isolated wall occurs, space seems to flow around it and a sense of anticipation is set up about the the hidden part of the space. It follows from this that a square space set up by walls that do not meet, although still providing a sense of enclosure, will be qualitatively different from one with entirely enclosing walls. Similarly if a ceiling appears to float above the walls, a different type of spatial experience is achieved.

As in the general case of layout or depth perception, textures play a significant role in our understanding of architectural space. Textural effects are not restricted just to the surface of materials; the effects arise at various scales. Thus, in a Gothic cathedral it is not only the quality of the stone that gives texture but also, at a different scale, the repetitive nature of tracery and decoration.

What does the look of a building say that it is? The railway station was a new building type in the 19th century and as such posed a problem to its designers—what should it resemble? To say that it should look like a railway station is tautological, of course. For the Victorians, the answer was, variously, gothic hötel-de-ville, castle, doric temple, etc. In all cases it was felt that something exceeding purely mechanical functionality was required. They were essentially theatrical spaces in which the drama of travel was acted out. Virtual Environments are perhaps our equivalents of the Victorian railway station: we have no direct historical precedents for designing them, and currently tend to fall back on the reproduction of other constructions which seem more or less appropriate.

Planning permission in Cyberspace

> In the instance of the monastery converted to a courthouse, who defined the first function? Who, and by what authority, transformed it? What physical changes were needed? Who named the functions?
>
> Markus 1993 p12

We have thought of Virtual Reality in terms of the solitary user, but already CSCW, MUDs, MOOs, Worlds Chat and other systems have made us realise the potential for shared computer 'space'. Many advantages are possible (see p85) but there are concomitant issues which are sure to arise sooner or later. We are currently only used to the model of the software publisher who sells a package offering certain functions, analogous perhaps to selling a product like a car, a power-drill or a food-processor. Virtual Environments and Virtual Objects which are made available for interaction by a single user perpetuate this familiar model. However, when we start to construct environments for shared use, the model cannot be sustained. Increasingly, in the physical world, people expect a say in the decisions made about their environment, and the same will be true of shared virtual worlds. Who decides what can be built? How are decisions taken? Some form of 'liberal democracy' will be needed to further majority decisions without squeezing out the needs of minorities. In an educational context, it is now considered good practice to allow learners to construct as well as to receive information—many courses offer their students the opportunity to engage in projects, as well as using the familiar techniques of written coursework and examinations. With this in mind, courses will also need to consider their attitude to VRML projects made by students.

Postmodern geography and the semantics of space

Geographers have taken an increasing interest in the semantics and the politics of space, whereas they traditionally saw themselves as dealing in the appreciation and mensuration of the physical landscape. To read any recent book on geography reveals a dramatic shift, which has been away from the 'obsession with landscape' (Jackson 1989) to such aspects as the significance of space in urban social theory; gender and the urban environment, posing a challenge to conventional distinctions between the public (spatial) and the private (personal); Orientalism (the concepts of the exotically distant); and the subjective geographies revealed by the characters of Hardy's Wessex novels (Jackson 1989 passim).

While it would be ridiculous to suggest that anyone designing Virtual Environments should feel obliged to read the vast literature of postmodern geography, an understanding of its general drift is useful, since once again it emphasises the meanings of spaces and how we use them.

What social functions converge in one building which were formerly separate? The supermarket now fulfils the functions of many small retail outlets. What functions which were once combined are separated? The pre-1834 poorhouse combined all those functions which subsequently devolved to the prison, hospital, school, asylum and workhouse.

Social views, political views even, inform the design of all spaces. Even gardens, assumed by many to be the epitome of detachment, have a sociopolitical history. The English landscape garden of the 18th Century was designed as a conscious rebuttal of the highly ordered, geometric gardens of the Continent and in particular France, seen as emblematic of authoritarian regimes by contrast with the freedom-loving, laissez-faire English. Pope mocked the whole idea of symmetry in a garden:

Grove nods at grove, each alley has a brother,And half the platform just reflects the other.
Pope, Epistle to Lord Burlington, NoIV, line 117

The history of architecture certainly seems to confirm that those buildings whose main aim is to impress authority, to inflate their own power in relation to the viewer, do indeed tend to be grandly symmetrical. To make a construction which others will move through and inhabit is to impose a decision on the user. It seems likely that the users of the future will not be happy to inhabit the consequences of someone else's remote decision-making.

Theatricalities—realism and non-realism

If it is accepted that VR can, and in many cases should, transcend the borders of physical realism and non-realism, an obvious place to seek inspiration is the theatre. As with painting, there has been a significant transition throughout this century from a view of the theatrical set as a realist pictorial view though a window, to the idea that the stage is an arena where the set-designer can do anything at all, take whatever freedoms s/he chooses, in the service of the drama. Ironically, this view of theatrical design has now itself been cited by Laurel (1991) as a model for interface design—a plea to interface designers to view their task as the construction of environments suitably equipped with 'props' in which the user's actions will take place.

The construction of an alternative reality as a context for drama might have something to teach us about the construction of virtual realities for education. Both seem to be about conveying meaning by constructing spaces which the viewer/user is invited to explore and interpret.

At the beginning of the 20th Century, the art of the theatrical designer was little theorised. There was a general assumption that, in addition to providing the entrances and exits, the furniture and the props, specified by the script, it was the job of the theatre designer to represent as accurately as possible the architecture, furniture and so forth of the period described in the play. That the set should do something other than represent the surfaces of a real scene was no considered. However there was, for example in the productions of Irving, an insistence on unity in the overall conception which already might begin to guide our thoughts about VR (Bablet 1962). In particular, contemporaries praised Irving's use of electric lighting to pull together the disparate elements of the scene (and no doubt also to hide in shadow some of the mechanical aspects of the set). As VR systems become more capable of handling complex rendering including the effect of multiple light sources, we shall almost certainly want to use lighting more selectively, and perhaps under the user's control, to make the VR 'speak' more effectively. A recent (non VR) project at the Centre for Electronic Arts made interesting use of darkness in which the user was enabled to shine a virtual torch, in an attempt to engage the user more fully in the scene.

Edward Gordon Craig (1872-1966) was a follower of Irving who took theatrical design in a direction of particular interest for this report, departing from the naturalism of his predecessors. Why did he do this? What was it about a wholly naturalistic set which he felt did not serve his purpose, and how might his thoughts guide us in our choices about virtual realities? Craig's reasons were based on a desire for what we might call media integration, concentration and consistency.

Media integration

By avoiding naturalism, the mimicry of reality, Craig found that he could better integrate all the elements of the production—in his case: actors, colour, music, movement (Bablet, 1962 p43). One of the problems of virtual environments is that we shall probably want to embed other media in them: for example, text. If we took the naturalistic route we would be obliged to put text into virtual books on virtual shelves, which seems to forfeit many of the advantages of non-physical media. To watch a movie, we would have to operate a virtual VCR. But by avoiding mimicry, we may be better able to create 'realities' which provide a seamless integration of the different media types, while at the same time keeping the strengths of each constituent medium.

Concentration

By being selective and non-naturalistic, Craig also found that he could focus attention where he wished. This raises an interesting question in relation to VR in Higher Education. On the one hand, we might be keen to create environments which, through non-realism like Craig's, focus users' attention on particular features, in a way that unfiltered reality makes it hard to do. On the other, we might want to create deliberately complex environments, where the onus is on the user to discover the truly important aspects, somewhat like the approach taken by (fictional) detectives when, faced with the chaos of the world, they are able to apply their attention to just those items which turn out to be important clues.

Also of importance to Craig, and included in the opening remarks of this report, was the question: is reality enough? One of Craig's problems with naturalism (at least, with unthinking naturalism), was its inadequacy to get beyond realism. There were several aspects to this. For example, in creating a scenario for Shakespeare, a naturalistic set could undermine the intentional non-realism of the words spoken (ibid, p45). Perhaps this seems an arcane problem, but in fact it is just a special case of the tension between symbolism and naturalism, and symbolism is by no means confined to the arts. When we consider something as simple as a graphical user interface, and consider how such a symbolic environment might be concretised as a virtual environment, it is clear that not all the advantages lie with a naturalistic solution. The non-natural symbolic nature of the GUI may actually be helpful, for example in making clear that this is only one of many possible metaphors for the operation of the computer, or in its ability to take liberties with scale so that all objects however near or distant have the same size.

Another Craigian idea valuable in thinking about virtual environments is the greater ability of non-naturalistic environments to suggest the unseen (ibid p134). In current practice, VR systems deal predominantly with surface appearances, and the systems in existence tend to emphasise the visual over the behavioural aspects of what they model—depiction rather than simulation. Craig's objection to naturalism in the theatre was partly based on his fear that it created a surface layer which distracts attention from the real messages of the play. The limitations of current VR technology have tended to enforce heavily simplified and stylised worlds upon us, so that we can not currently begin to create virtual worlds containing all the rich minutiae of the real world. A Craigian reduction to elementals has been forced on us by the technology! But when the naturalistic detail of the real world becomes more achievable in VR, we shall need to decide whether this is, for some purposes, both more and less than what we want: more because details may distract us from what is being said; less because we may need powerful symbols, and symbols normally work precisely by eliminating particularities of surface appearance. A historical example might be the work of the pre-Raphaelite painters, whose work seem now to lack any clarity of meaning but which preserves in minute detail the (irrelevant) surface characteristics of what the painter observed. Recent (non VR) projects in the Centre for Electronic Arts have tried to grapple with the problem of stripping away surface excrescence from multimedia products, resolving complex functionality into simple interfaces, while commercial software development by contrast has tended to grow extra buttons to deal with extra functions, so that the user feels s/he is interfacing with a layer of control devices, rather than with information itself. Reality is full of clutter, and we do not necessarily want our VRs to be the same.

Consistency

We remarked that the theatre has been adopted by Laurel as a metaphor for the human interface. One of the points on which she and Craig concur is in emphasising that the stage, considered as a static environment, is not in itself the point, that it is a location for action. Craig was untypical of his time in emphasising 'action—movement—dance' rather than speech, literature or poetry in the theatre. Clearly, part of the question 'What can we use VR for in HE?' must include, 'What should happen in these Virtual Realities?' We shall need to choose whether our VRs include only action by a single protagonist, the user; by multiple users, probably networked; whether the VR is essentially inert, responsive to the user's actions or whether it includes objects and agents with autonomous behaviour. This in turn means considering how people interact in shared spaces (see p35, 83), as well as whether the user has a first person or third-person relationship to the environment (see p55). Strangely, Craig became increasingly fascinated by the possibilities of a purely mechanical performance without the irritating distraction of real actors! He began to feel that sophisticated marionettes would be preferable (ibid p107 passim). He would surely have loved VR.

For Craig, the emphasis on the set as a site for action led to his facilitating in his designs the process of change, often by the re-configuration of modules. He likened the set to a face, which always has the same parts laid out in a recognisable pattern, but which alters its expression with any change in one of its features (Bablet 1962, p 126). Given the potential flexibility and reconfigurability of computer models, this is something we may want to consider in making virtual environments, potentially achieving both the unity and security which a single virtual 'place' provides, at the same time as the flexibility and multiple functionality of many places. Commercially this will be important, since it is a way in which an identifiable 'brand' of place can have many uses. While real buildings and landscapes are not generally reconfigurable, virtual environments will surely want to reject adherence to this limitation of the real world.

It might be thought that this report is advocating an anti-naturalist approach to VR, such as Craig adopted in the theatre, but this is not the case. The point being made is that in deciding on the role of naturalism in Virtual Environments, some of the reasons that Craig had for (in his case) objecting to naturalism may be worth considering for our own purposes. Above all Craig was a functionalist, in that he saw the purpose of the design as to serve the idea (not the other way round).

Omission and abstraction

Craig emphasised the value of omission and abstraction at the expense of naturalism. It is interesting to look at film and see what it omits, while at the same time seeming to be a total experience, whether for purposes of dramatic fiction or factual communication. Notably, it relies on only two sensory channels—vision and sound. There have been attempts to incorporate smell in cinema since at least 1906, with Smell-O-Vision and Aroma-Rama both being launched in the late 1950s. Neither system proved popular, although technically they were both successful: the idea of scenting films was abandoned (Katz 1994). This raises the question of what other economies we can make in our virtual worlds.

In 1851 a Stereoscope was exhibited in the Great Exhibition, and was patronised by Albert and Victoria. A viewing device with two lenses, it enabled a pair of photographs made with two adjacent camera lenses to record stereoscopic views. By the mid 1850s, one version had sold more than a million in England alone, and by the end of the decade the ambition of the London Stereoscope Company, 'no home without a stereoscope', was almost fulfilled (Macdonald 1979). All manner of exotica were available—ghost pictures, moral tableaux, freaks and oddities. There was also pornography. Today the stereograph is relegated to the status of a child's toy, not taken seriously as a medium of communication. Crucially, its ability to present the third dimension is seen as unnecessary. What might have led to its demise, while straightforward photography has continued to increase every year since its invention? Stereoscopes were available quite cheaply, and many were originally bought, but the craze passed. Perhaps a means of communication which required a mechanical device for viewing could not be assimilated into normal life? But what about the gramophone and the television, both dependent on physical machinery? In the end, as designers, we are always free to make decisions about what can and should be omitted, and it is surprising what effective use can be made of media which offer only one or two channels of address.

The following quotation may cast an additional historical light on this issue:

> There seems to be little doubt that there will be colour. When colour is perfected technically, the film will be an almost perfect copy of reality, almost 'the real thing.' With the perfection of the magnascope, the screen will be made to issue into the auditorium: the audience will be enveloped in the film. The approximation to reality will be still more convincing, the opportunity of imaginative creative work still more improbable, the cheap appeals to participate in a crude and unhealthy experience still more compelling.
>
> William Hunter 1932 Scrutiny of Cinema p11

More spaces in literary imagination

The accepted definition of virtual reality of course relates to computer-stored, primarily visual, models which achieve their aims by visual specification of the attributes of objects and environments. We have looked at analogies with architecture and the theatre. It is worth thinking about the completely contrasted approach taken by books, especially novels, which achieve their effects with a bare minimum of specification. There is much less data in a book than any but the simplest computer graphic models, yet most readers would probably agree that the paucity of data in a book is not in itself a problem—that in fact books acquire much of their effectiveness through omission.

The gamut of literature covers the whole range of relationships between representation and reality, from the highly realistic to the fantastical. Literature provides a rich resource of ideas about space and spatiality, which we suspect is largely unresearched (exceptions include Barrell 1982). The essential idea is again to avoid assuming that space just 'is what it is', an inert and essentially irrelevant location for the action, and instead to arm ourselves with ideas about spaces with meanings, spaces which facilitate and augment the action as much as do the personalities of the characters.

Barriers
We have already commented on art and literature which depicts stratified societies, and emphasises the ease or difficulty of moving from one layer to another. A similar device though used for different purposes is proposed by Lewis Carroll in Alice through the Looking Glass. Once again we have a model based on some characteristics of real spaces, but with all the important qualities of the vision arising from the tension between this realism and the elements of unreality.

Pilgrimages
Another class of interesting proto-virtual spaces in literature, is where travel along a route is put forward as an analogue of spiritual progress. Again there is a long tradition: it is only necessary to mention Pilgrim's Progress and the Yellow Brick Road! In fact these linear spaces are a special case of a whole category of constructions placing symbolic elements in spaces.

From space to place Distinct, but not separate, from the concepts of space, is that of place—the particularities of a location or a type of location—what differentiates a cathedral from a railway station, or this particular library from that library.

> All really inhabited space bears the essence of the notion of home.
>
> Bachelard 1958

One of the most influential writers in this regard is Bachelard, whose 1958 book, the Poetics of Space, includes among others chapters entitled House and Universe, Shells, Corners, and The Dialectics of Outside and Inside. The great value of the book lies in its extensive coverage not of spatiality in general, but of the characteristics of particular kinds of spaces. While spatiality may be partly perceived in an a-cultural way, as we said before, clearly the characteristics of particular kinds of spaces are heavy-laden with cultural influences. For a space to be perceived as a house, makes a highly influential difference to the way it is 'read', compared with it being read as a public institution. In VR, where we have freedom to generate new space-types which evoke any existing genres of space, plus eventually completely new ones, what a space is perceived to be matters enormously. Both the thoughts and behaviours of users, their psychology and their actions, will be heavily influenced by whether a space is perceived as of the type normally thought of as public or private, intimate or challenging, large or small.

Large and small

Another virtue of Bachelard's work is that it deals with scale, explicitly in the chapter Miniature, but also as a theme throughout the book. Most discussions of spatial constructions concentrate on spaces of a given size, such as cities, public buildings, jewel-caskets. It is unusual to find the issue of scale raised at all. Scale is an issue for VR especially of the non-immersive kind, because of VR's disembodied nature. In reality, our physicality prevents us from reaching up to things higher than ourselves and from climbing inside things smaller than us. VR knows no such prohibitions, unless they are deliberately constructed. This is yet another way in which VR cannot, by its nature, replicate the lived experience of reality, and must remain always at some level a representation.

Virtuality

> [The body] works in Euclidean space, but it only works there. It sees in a projective space; it touches, caresses and feels in a topological space; it suffers in another; hears and communicates in a third; and so forth...
>
> Serres 1977

To an extent we can consider spatiality as something merely perceived, without any action by the user, but the same is not true of virtuality, which is essentially about the relationship between the user and the system.

Virtuality builds on the idea of Direct Manipulation popularised by Shneiderman (1992), discussed below. It aims to produce forms of human computer interaction which give users the feeling that they are engaging with data rather than with tools to manipulate data. Hutchins et al (1986) describe 'the feeling of involvement directly with the world of objects rather than of communicating with an intermediary.' Instead of turning on a light by touching a light-switch, we simply touch the light itself. Virtuality is opposed to the idea of the Separable Interface (Edmonds 1992) which is applied as a control surface after the functionality has been decided on.

Nelson (1990) gives as an example of effective virtuality a terrestrial globe: 'A globe does not say "Good Morning"; it does not bother you with menus, icons or prompts. You turn it and move your head to the most useful position for overview or detail, that's all.' A virtual globe is indeed an interesting case for any discussion of VR. It has the great advantage of eliminating one kind of representation—that of mapping the world onto a 2D surface with all the unhelpful distortion that implies. We might imagine a time when students never again see a 2D mapping of the world. Nevertheless it is obvious from what we know of maps that many issues raised by 2D map-making are NOT eliminated by making a 3D virtual globe. For example when the user first encounters the globe, which country is centred in their view? Which way is up? —there is no scientific reason to put North at the top of a globe. We must presumably have the freedom to change from physical to political and other models. And as with all direct translations of physical dimensions into computer dimensions, a virtual globe of course privileges a particular form of measuring, whereas 2D maps have been able to exploit the freedom of their greater separation from actuality by, for example, using the depicted size of countries to represent their GDP, rather than their physical size. So Nelson's example of virtuality is not quite as persuasive as he makes out.

Direct manipulation and virtuality

The classical definition of Direct manipulation requires:

- continuous representation of the objects and action of interest
- physical actions or presses of labelled buttons instead of complex syntax
- rapid incremental reversible operations whose effect on the object of interest is immediately visible.

Shneiderman's thoughts on Direct Manipulation (1992 p204) are germane to any discussion of Virtual Reality. In particular he lists some difficulties with visual and spatial representations compared with text:

- graphical representations may be more extravagant in their demand on screen-space than the corresponding text [and of course on computing power]

- the low density of information in pictorial layouts may inhibit performance of tasks compared with, say well designed textual or numeric tables
- if visuals are used iconically, they must be learned [this is another factor in the debate between mimicry in VR and the construction of semi- or wholly abstract 'information landscapes']
- the visual representation may be misleading, causing users to predict behaviours in the system which it does not in fact permit
- choosing the right objects and actions is not an easy task

While Shneiderman focuses on relatively simple graphical interfaces, these points serve yet again to emphasise that the construction of Virtual Environments is an intentional, design-based activity which must take into account many factors, not a mechanical process of developing electronic counterparts of slices of the world.

However, Shneiderman does not mention what is perhaps one of the most problematic aspects of both Direct Manipulation interfaces and VR, namely that in abjuring abstract, symbolic notations we may suffer an intellectual loss as well as some gain. In tying ourselves to a visual, physical paradigm we may limit what we can think. A simple example would be that of a child learning mathematics. First the child does simple arithmetic by manipulating beads; later the child learns that number can be represented by abstract symbols; later still, realisation dawns that mathematics which cannot be done with the physical beads can be done with the symbol system. For example, zero and negative numbers are both concepts which are inaccessible without a symbolic notation, and both are indispensable to anything beyond elementary mathematics.

This is a problem of concretisation which needs to be acknowledged in VR generally, but especially in relation to Education. It could be argued that some of our greatest intellectual achievements arise from the mental manipulation of symbol systems, most especially text and mathematics, and it would be an odd outcome of adopting concretising approaches, via virtual reality, to deny ourselves and our students the ability to think in terms other than the physical. After all, most of our current education system is arguably geared toward encouraging the facility of symbolic thought.

Virtuality and human error

At first sight, it might seem that in constructing virtual worlds we can escape from the problems of human error associated with more abstract computer interfaces. However, no system can be devised which is not susceptible to human error: errors of conception—'mistakes'—and errors of action— 'slips'—will both occur (Norman 1988). Nevertheless, the kinds of errors are likely to be different. With command-line interfaces (CLIs) and with graphical user interfaces (GUIs), slips are all too easy for a user who is manipulating objects. For example, when deleting files, the user of a CLI may incorrectly type a command, irretrievably deleting the wrong files; in a GUI, where the user may accidentally drag the wrong objects into the trash, there will normally be some limited possibility of recovery (but only if the trash has not immediately been emptied by the user); in a virtual environment, we might hope that the greater variety of visual cues possible in a Virtual Environment will militate against this sort of slip. Colour, size, texture, transparency, position, orientation, and above all resemblance to some real-world referent (for example objects belonging to individuals might bear their signature, portrait or other characteristic marker), can all invoke our associative skills to correctly identify objects. However, any user of a 3D modeller will be familiar with problems of virtual objects and environments which are not so easily solved, for examples confusions about scale (VRML 1.0 used an arbitrary measurement unit which predisposed designers to give insufficient consideration to the ‘actual’ size of what they were building), and about whether the model or the user is moving (a problem of desktop VR rather than immersive VR).

Design faults in virtual environments

In addition to those new problems associated with virtuality, we may expect to transfer some of the problems we have with 3D environments in the physical world to the computer domain. Some problems arising from poorly designed environments and objects include:

- inadequate identification and signage
- lack of affordances
- poor use of mappings
- inadequate feedback

It is assumed that virtual environments will be interactive in the full sense, so that the users' interaction is not confined to altering viewpoint. VRML is changing to accommodate such needs.

Inadequate identification and signage
As in large buildings and cities, poorly designed Virtual Environments will have inadequate identifying features and insufficient navigational cues (Passini 1992).

Lack of affordances
Affordances are those traits of a designed artefact that 'tell' us how to use it (based on Gibson's ideas about visual perception (1979) and used by Norman (1988) and others). A well designed handle will look like an object that can be turned, while on the other hand the multiplicity of identical buttons on a consumer electronics device normally give no clue—other than their labelling—to their purpose or mode of operation. It would be sad if we built a virtual VCR as appalling as most of the actual VCRs on the market (Thimbleby 1991). There is some reason to believe that the visual uniformity and resultant incomprehensibility of electronic devices arises from the freedom from physical constraints in their construction. If building a steam-engine out of brass and steel, there are many limitations of physics on what can be designed; but if building a music amplifier, the physical components (being small and not easily identified by the lay person) cannot either 'explain themselves' nor dictate the form of the assembled article—another black rectangular box is produced. The relevance of this to virtual objects of course is, that an even greater disconnection is possible for a virtual artefact between its form and its behaviour or mode of use. The designer of virtual objects has the freedom to make even worse objects than those commonly available in the real world. However by applying a proper understanding of human perception and behaviour, and transferring to virtuality the best of design principles from the familiar world, the designer also has the freedom to exceed the usability of real objects. This is design considered as a gift to the user, rather than as an opportunity for the designer's self-indulgence.

Poor use of mappings
Other aspects of physicality have influenced the form of real-world artefacts. The form of objects and constructed environments is not only determined by the nature of their materials but by the fact that they will interface with human beings. A door-knob is the right size and shape to fit the hand. If the mapping of virtual objects to the user were well designed, then the forms designed for use immersively with a dataglove would be dramatically different from those designed for desktop use with a mouse-pointer.

Inadequate feedback
The objects with which we interact must be capable of registering our actions. Feedback should be clear, instant and appropriate. GUIs have made considerable strides in providing kinds of feedback, including:

- continuous representation of the object of interest—a defining characteristic of DirectManipulation—such as the file being moved across the desktop
- input-acknowledge indicators, such as buttons which highlight when the pointer intersects them
- progress indicators, such as a rotating cursor to indicate a delay while processing takes place
- error messages of many kinds

In current VRML browsers, it is easy to be unsure whether a delay is, for example, caused by a failure of user-input such as a mouse-click or by processing delays. Just as the language of feedback in current GUIs has evolved over a period of more than ten years, it can be assumed that it will take a while for an appropriate range of feedback styles to mature in Virtual Environments: another difficult but stimulating task facing the designers of Virtual Realities.

What you see is ALL you get

The history of desktop computing is sometimes taken to demonstrate that GUIs are superior in every respect to CLIs. A single example of poor design in many GUIs will demonstrate that this is not so, and point to problems which we must avoid in Virtual Environments. Using a CLI, we can interact with objects regardless of whether they are currently represented on screen (for example copy a:myfile.txt b: may be used even when neither file's name was previously on the display). By extension, we can also easily set up batch processes to work on a set of objects over a period, to which we need not attend. But in some GUIs, we can interact only with those objects currently in view, and we must manually undertake each interaction. The significance of this is that it is symptomatic of the excessive concretising which we have warned against elsewhere in this report: it is an unfortunate corollary of elevating the concrete at the expense of the symbolic. The fact that in a Virtual Environment we can interact with Virtual Objects by direct manipulation, should not mean that we must do so.

More prerequisites for successful virtuality

What are the preconditions for successful virtuality? What prompts the sense of presence, of directness, of engagement? Part of the point of VR systems is their very sense of 'realness' (which we would argue is not dependent on mimicry of real world referents), and with sufficient computational power and the design of effective feedback already mentioned, we can also hope for engagement to arise from the way in which objects in the scene will react to the user's actions. With current technology, a responsive 2D system may be more satisfying than an unresponsive 3D one.

In addition to these obvious requirements, there may be more subtle psychological characteristics which would assist. The ideas of Csikszentmihalyi (1988, Bulmer 1995) offer an additional insight into what we might mean by successful virtuality. He has tried over an extended period to gather evidence of what it means for a person to 'lose themselves' in an activity, to achieve a state of experience which he has named 'Flow', whose characteristics include:

- a sense of complete engagement
- alteration of self-perception (both less and more perception of self)
- increased perception of external systems and devices as extensions of self

One of his useful observations is that this kind of super-engagement arises intrinsically from the activity. While an activity may be begun for extrinsic reasons, it later becomes self-motivating. This seems to apply equally to work activities and recreation.

Csikszentmihalyi is less useful when it comes to helping us state what might be the preconditions for Flow. He suggests that the difficulty of the task in hand should be maintained between the boringly easy and the frustratingly difficult, hardly providing new insights unknown to the Higher Education community! His populist writing is aimed less at the designers of systems than with the attitude of the participant. Nevertheless, his work may be worth study for anyone trying to devise inherently absorbing systems. We must always remind ourselves that the motivation which users may experience when first they encounter Virtual Environments will be transient, assuming that VR will become a commonplace mode of interaction.

Academics, and especially the creators of courseware, sometimes look enviously on the engagement which computer games bring to users, and games are certainly among the first widespread uses of 3D interaction. There seems little doubt that the provision of intriguing, puzzling, often confrontational environments assists in engaging the user. However, we must remember that an important contribution to the appeal of games is their very lack of 'content' and that they are elective—people are not obliged to play a computer game in the way that they are obliged to take a particular course module. Some of the attraction of computer games is therefore nothing to do with their design, and Education can never by its nature achieve the same style of devotion.

Environment into process

Laurillard (1993) has argued that hypertext is not an educational medium, in the sense that it provides only an environment in which learning may take place, but falls short of helping form the learner's actions. Likewise a library is not a course. In a sense the educator must stand in the role of narrator, leading the learner through experiences: even in intensively resource-based learning, this role for the academic survives. A similar objection could in principle be raised to Virtual Environments, and most designers of Virtual Environments assume that other media will be needed to assist the learner in getting value out of the environments and objects offered.

However, a special case of the affordances already mentioned is worth considering in this context. Some kinds of environments seem to invite the user to explore them in a particular way. When people describe well-designed gardens or buildings, it is common for them to read into these passive environments a formative, perhaps narrative, quality to the space. One example stands for many:

> The resulting spatial rhythm is much smoother than that of Romanesque cathedrals or of Noyon. It is no longer split into numerous units which one has to add up mentally, as it were, to summarise the spatial totality, but concentrated in a few, in fact three, sections: west, centre, east. [...] It leads you towards the altar as forcibly as did the columns of Early Christian basilicas.
>
> Pevsner, 1943

Clearly the architecture invited certain trajectories on the part of the user. A recent project in building a VRML environment, at the Centre for Electronic Arts, explored the idea of concealment as a motivator—users tended to look into a space which they could glimpse but not wholly see. Again we are extending concepts familiar from recent attitudes to interface design: the designer must work not on the artefact itself, without taking into account the active processes of the user: what the user will see, hear, and above all do, in the environment.

Protagonists and point-of-view—a case study from film

An obvious alternative to the idea of spaces which to some extent 'explain themselves' is that of inserting either a narrator or a protagonist to assist the user. Though confined to the presentation of fiction, feature film offers some thought-provoking exemplars, among them 'The Lady in the Lake' directed by Robert Montgomery and released by MGM in 1946.

In this film some of the conventions of Hollywood film-making are overthrown, especially in relation to the viewer's relationship to the protagonist. Through our familiarity with film we have learned to 'look through' the conventional structure narrative devices so that we hardly notice them, but they are conventional and not 'natural'. For example, who is it who sees what is happening in a film? Is it the viewer or the narrator? When the director wishes to show us, not what the protagonist is doing, but what the protagonist sees, then there are well-established codes for making this clear (typically a close-up of the eyes of the protagonist, or a shot which shows the direction of view of the protagonist in relation to other actors). Despite the fact that we are watching all the characters from the 'outside', we often 'identify' (or at least empathise) with the main protagonist. This effect can be strongly reinforced by the use of a 'first-person' voice-over narration. In the Lady in the Lake, however, we hardly ever see the protagonist, for the simple reason that the whole film is shot as it were through the eyes of the main character—he only appears in mirrors.

Why is it worth looking at this film, which is after all an authorial, narrative piece of fiction? Attempting to evaluate which aspects of the technique seem to work, and which not, may help us to foresee how in a realistic, filmic Virtual Environment we could best locate the user, and what problems the user may have in relating to this 'other world'. In evaluating strengths and weaknesses however, we have to acknowledge the prejudicial effect of more familiar, established filmic codes. What we see as 'right' will probably change.

Aspects of the film which seem unsuccessful (some of which have direct counterparts in existing VR techniques) include:

- A quick glance from one part of a scene to another is never used. The camera must pan cumbrously in a way quite unlike human vision. Ironically, jump-cuts (where one scene simply stops and another start) may be more analogous to how we use our eyes in the physical world.
- Likewise, when our attention is summoned from behind, it is difficult to convey the swiftness with which we can turn our attention to the source.
- Parts of 'our' body must intrude in the scene. For example 'we' strike a man with our fist; we turn a door-knob with our hand. Of course this hand is unlikely to resemble the viewer's own.
- Looking in a mirror reveals our 'self' which again does not resemble the viewer. Perhaps experience with 'avatars' (representations which 'stand-in' for a person) in shared virtual environments suggests that our problems with this 'other self' will be overcome (see Social space and Virtual Worlds, next) ?

Some aspects of this first-person/third-person view which seem more successful include:

- The sense of immediacy when we eaves-drop on another's conversation—we have a sense of really being there, hearing but unheard, seeing but unseen.
- When characters address us directly or respond directly to what we do (such as showing fear) the effect is more potent than in conventional film.
- Similarly, we feel guilty when we seem to hurt another character's feelings!

- Direct assaults on 'us' are perceived more viscerally, less intellectually than when observed from a remote standpoint. This effect of course is widely used by violent computer games, some of which allow the player to switch viewpoint at any time between being in the action and observing it from without.

One of the points of this extended discussion, apart from simply drawing attention to an aberrant attempt at virtuality within the Hollywood tradition, is to recall that the filmic conventions which we find so convincing and unremarkable in the cinema themselves took time to evolve. Again we can only wonder what new forms of expression may be latent within VR. The early cinema was confined to simple documentary sequences and technical showing-off of various kinds: no one could have predicted the major role of the cinema in our culture.

Social space and Virtual Worlds

The attention of software designers has shifted in recent years to supporting shared workspaces and collaborative working of various kinds. At the same time in some Higher Education institutions there has been increased interest in students undertaking collaborative work, which is seen both as assisting in the educational process and in improving the student's preparation for working life. A partial shift to assessment by project-based coursework has also contributed to changing attitudes.

While some effective collaborative environments may not be actually visualised at all, others are given the visual attributes of actual spaces using VR techniques. These worlds can exploit social aspects to encourage the sense of place, rather than using purely perceptual cues to suggest only space. Shared worlds are made possible by changes to VRML. The VRML-2 standard aims not to inherently provide the new means to do this—it is not part of the standard itself—but instead to provide handles for the functionality to be added by additional routines.

Issues arise for such worlds as to how users—represented by avatars—can engage in this space. Currently, most Internet-based, VRML-oriented, multi-user worlds, such as CyberGate, allow the user to be represented with a standard avatar image. More recently it has been possible to create one's own avatar with VRML. However, these inevitably look somewhat cartoon-like. 'Social problems' have occurred previously in text-based shared worlds which may emerge also in the new 3-D worlds as these become more widespread. Having been available far longer, the text-based worlds have exhibited the problems first.

In particular there is a problem in the way that both the spaces and the other avatars are perceived by users. For some they are pure fictions, based as the technology is on games, and this characteristic is likely to be exacerbated by the cartoon-like graphics of the 3-D worlds. Other users, on the other hand, see their avatar as an extension of themselves. This mismatch has lead to many conflicting situations, such as the now infamous 'rape in Cyberspace', in which one user knew the system sufficiently well to manipulate the words and actions of another's avatar.

One approach to addressing this issue is to ground the avatar more firmly with the user. This is possible for example through the use of the user's own voice, rather than the use of typed text. The avatar can be based on one's photograph, rather than a primitive VRML model. And in future, it will presumably be possible to stream a video of one's face onto the avatar.

Augmented Spaces

Another approach to alleviate such problems could be to 'de-virtualise' the world by connecting it to a real space. Recent work brings the value of complete 'virtualness' into question, and a number of projects are concerned with linking the virtual to the real world to create 'augmented realities' in which digital technologies do not replace, but supplement, that reality. An example of such ideas is the work at Lancaster University (http://www.lancs.ac.uk/computing/users/andy/vrml/skylabspy.wrl). In this world, a 3D model of the research lab is connected to the real-world lab, and when people are logged into machines, their avatars appear in the virtual world standing before a model of a terminal. It is then possible to phone them, or use some synchronous computer communications, so establishing real-time contact.

A similar project is being undertaken at the University of Nottingham, who are building an Internet Foyer (http://www.crg.cs.nott.ac.uk/Foyer/foyer.html). This will link a three-dimensional model, based on their web site, to the physical foyer of the building. The model is projected onto the wall of the foyer, and the foyer is visible, through a camera, to the remote users. The intention is that users visiting either the web site or the foyer are aware of each other's presence: extending the social function of the foyer-space to encompass the virtual.

It is possible that by connecting the virtual world into a real space, the participant will become successfully 'grounded' in it. That is, users will feel that they are visiting a space in which certain types of behaviour are acceptable and others not. For example, if someone feels that they are, in some sense, in a foyer within Nottingham University, then they are less likely to treat the experience as a fiction or a game, and act accordingly.

Locales

The social nature of these shared worlds also affects the way they are constructed spatially; indeed, the value of the spatial metaphor becomes more apparent. One common approach is to divide the space into areas, usually called locales, that allow for the subdivision of activity. Even non-graphical virtual worlds, such as the text-based MOOs, use this metaphor most successfully. Giddens (1984) offers us a means of understanding this with his interpretation of locales—areas that set the scene for certain actions. Within locales, the affordances of the environment, together with the nature of the tools provided, predisposes each locale to particular types of work and interaction, by analogy with the way in which in the real world we have, for example, offices, coffee rooms, and foyers. The spatial relation of such a foyer with other spaces is important: it acts as an entry point to other more private spaces.

A number of collaborative tools have drawn on this idea, such as the Worlds project (Fitzpatrick, Tolone & Kaplan 1995).

Representation

We looked previously at the issue of what is included in the model and what is omitted, but we need now to turn our attention to how that model is represented visually.

Issues of realism and photorealism

Our most obvious inheritance in terms of representation comprises the two traditions of painting and photography. It is useful to compare them. At the end of the nineteenth century, a view emerged that the history of painting was one of perfecting the imitation of reality. Ruskin instructed painters to “Paint all you see, selecting nothing, rejecting nothing.” However, the subsequent history of painting has drawn attention to the wide variety of representations available, few of them photographic (even photo-real paintings have always filtered and configured the view). In fact the history of our understanding of perception demonstrates that we do not know what we see. For example, it seems likely that the design of early cinematography lenses was based on a view that the human eye sees everything in focus at once—a view that no one now holds. A different opinion about that perceptual process would have led to a different concept of photorealism.

The development of rendering algorithms has largely been devoted to achieving a resemblance to the sorts of surfaces depicted by photography. Lansdown and Schofield (1995) however list some of the attempts that have been made in other directions, including Schofield’s own Piranesi renderer. They document the attempts which have been made in both 3D rendering and paint systems to replicate the materials traditionally used by artists, but point out that this is in its own way a limiting view. What is required, they argue, is not just the facility to imitate familiar rendering styles, but techniques to make more expressive representations. They do not use this primarily to refer to emotional expressiveness, but to the ability that some drawings, for example, have to ‘speak’ to us more directly of the forms they depict than could a photograph. One of the most profound skills of drawing is the use of mark-making as a form of explanation—the marks take on the status of a form of language, helping the user to ‘read’ the forms depicted. This is why drawings, paintings and sketches often contain marks which do not directly correspond to any observed edge or tonal difference in the scene but instead, for example, establish the spatial relationships of objects to one another, or the dominant directions of a surface. In fact, many kinds of drawing freely mix marks with obvious direct referents in the observed scene, such as the edges of objects, with other marks of more schematic kinds. The person skilled in drawing is able to combine in one artefact the traditions of verisimilitude to the retinal image and those of schematic representation.

Perspective systems to choose from

> [Painters] were enabled to ‘capture’ aspects of that world in a wholly unprecedented form of synthesis, which owed as much to science and mathematics as it did to the earlier practices of pictorial art.
>
> Bann 1987

We cannot deal here with the vexed question of whether three-point perspective (the one we normally think of as ‘perspective’) is in a special sense, correct. However, we can look at the status of ‘perspective’ and ask ourselves whether it is the only system we would want to use to project our virtual realities.

When three-dimensional cartesian data is stored in a computer system, this implies almost nothing about what the user will see. Assuming that the eyes are the user's means of sensing the three-dimensional model, then it is unavoidable that the three-dimensional data will be mapped onto two-dimensional surfaces, whether a conventional desktop monitor, projection screen, miniature LCDs in a headset, or even the retinae themselves. A 3D to 2D transformation must be applied, and what this transformation may be is dictated by the rendering system, not by the model.

In a graphics rendering system, any consistent model of perspective can be applied (it could even be inconsistent if we chose to make it so). It is a simple matter to render cartesian three-space in any of the standard projection systems: orthographic, axonometric, isometric, two-point or three-point perspective. However, VR systems are normally assumed to map the three-dimensional data using three-point perspective. Arguably we are back to a question of definition, since some would certainly want to define VR as based on three-point perspective projection. We would argue however that the essential qualities of spatiality and virtuality are present even when other perspective systems are chosen. In addition as we shall briefly see, even the adoption of 3-point perspective leaves many questions undecided.

A theme of this report is the recurrent question, 'Does it work?' In evaluating the role of VR in Higher Education, we have chosen to look above all at effectiveness. Therefore we must ask for a given projection system, not 'Is it correct?' but 'Does it serve our purpose?' Clearly there will not be a right answer for all cases.

What is three-point perspective for?

> The relationship between what we see and what we know is never settled.
>
> Berger 1972

One point to note immediately is that in the mass of the world's visual representations, 3-point perspective is unusual. In the past it has been supposed that the acquisition of knowledge about how to use these systems is symptomatic of the essential superiority of post-Renaissance European culture over all others, and there are probably residues of this opinion current now. The core of such an idea is that other cultures would have made use of three-point perspective if only they had been fortunate or clever enough to invent it.

There has been much speculation about why the system was invented/ discovered (even the choice of word is loaded) at the time and place it was, and arguments have included chance, the rise of the property-owning class, a shift away from a theocentric view of creation, and many more.

Because we are brought up with 3-point perspective, it seems natural—not just the obvious choice, but the one we would instinctively choose. In some senses however, 3-point perspective seems not to be natural at all, and perhaps this is why it is so uncommon as a form of representation. Teaching drawing (in the West) involves inculcating the process of capturing on paper the three-dimensional world from a single viewpoint according to the rules of 3-point perspective, and it turns out to be an extraordinarily difficult to do. Some learners never manage the essential skill of 'seeing' the top of a rectangular table as a trapezoid at a particular angle to the line of sight: they can only see what they 'know is there', a viewpoint-less view. Following an ecological line of argument, it is interesting to speculate why it might be that (1) seeing in perspective is an unnatural skill acquired by most people only with substantial difficulty and (2) it is very uncommon in the history of the world's cultures. We might begin to suspect that, from an ecological point of view, it is somehow 'unnecessary', or 'unproductive'.

Until the invention of photography (which in its usual form produces images in three-point perspective), the application of 3-point perspective to images was almost entirely in the field of Fine Art painting. Even after photography, there has been a continuous outpouring of images which do not use 3-point perspective. The representations favoured by architecture, product design, engineering, magazine advertising and many other image-using trades and professions use 3-point perspective as the exception rather than the rule. This is not because of a lack of technical capability, but because in more cases than not the 3-point perspective image would not serve the valuable purpose that (for example) an orthographic image does. It would not be capable of conveying the necessary information. It is too easy to assume that 3-point perspective is a sort of superset of projections which contains more than each of the lesser projections individually, but of course the ambiguities of 3-point perspective can often only be avoided by adopting some simpler projection altogether.

What value do we get out of a 3-point perspective image? We might be able to answer this question by separating images into those generated by photography and those constructed by other means:

- Photographic images (photography, video, film) are valued above all for their truth about the past; drawn 3-point perspective images for their concreteness in displaying the future.
- Drawing, to a much greater extent than photography, is valued for its ability to select, to abstract.
- Some forms of drawing can also capture the dynamics of the process that created them.
- Some forms of drawing, especially those called 'sketching', can also be productively ambiguous.

The truth in photography

> The flat, usually rectangular, images which are photographs make a claim to be true that painting can never make.
>
> Sontag 1977 p86

Photographic images have a reputation for showing things as they really are. This is not the place to debate the truth or otherwise of photography, nor to dwell on the practices which can assist photography in being untruthful (Sontag p115 passim is useful on the various claims made for photography's relationship to knowledge). Many have stood in front of a painting of a historical figure and thought, I wonder what s/he really looked like?' and have had little doubt that a photograph would have helped to answer this question. Whatever the truth or otherwise of photography, when it comes to belief, it is unrivalled. Academics would encourage a student to treat with some scepticism a photograph from, say, an electron microscope, or from a historical archive of 1930s Britain, but not the same kind of scepticism that we would adopt towards a drawing. We believe in what photography shows us.

It is interesting to speculate what the 'belief value' of VR will be in future. While VRs remain clearly constructions, they are likely to be regarded as low in believability, however realistic (in another sense) they may look. But if a practical means can be made available whereby VR models can be created as a counterpart to reality (whether using multiple cameras, laser-scanning or whatever), in the same way that photographs are perceived as an unmediated response to the light-values of a scene, then the perceived authenticity of such models will receive a major fillip.

Of course the question of authorship and the status of the author will also be influential: the triad - authorship, authority, authenticity - will continue to be an issue as it is in relation to any medium.

Interestingly, the belief we have in photography seems to illicitly carry over into our appreciation of overtly fictional works, so that films are seen as in quite a significant way more real' than plays (despite the fact that in a play the actors are really there).

The truth in drawing

By drawing' here we mean any non-photographic image construction. Perspective drawing allows us to create convincing' representations of the future, and of things which are unseeable (such as atomic structures). This is one of the aspects of VR - which in this sense we can regard as a form of perspective drawing - which has excited most interest. Even when we forgo the realism about things seen which photography seems to offer, the feeling of truth to appearances which 3-point perspective drawings provide is one of their greatest advantages.

Selection in drawing and photography

Another important aspect of the relationship between the model stored and the scene rendered is that of selection. In photography or film/video, selection is achieved by a number of means, but especially by framing - deciding at what to point the camera. This is an authorial prerogative which we have deliberately denied ourselves in VR, where we normally allow users to look where they please. However, the issue of selection still arises. though in a different form.

Some VR systems may resemble photography in the respect that whatever appears in the environment (the stored model) will be seen (rendered). If the mapping of a complex 3-space construction onto the visual display is done wholesale, without any ability to suppress detail, to render some objects translucently or omit them altogether, then the user may be oppressed by irrelevant data. To avoid this, we will again want to take liberties with our virtual realities, deciding what to show and what to suppress, in essence, filtering' the model. Levels of detail go a small way toward implementing this concept.

One of the reasons why drawing and other forms of constructed' representation still have a role since photography is that they allow constructive filtering. The technical illustrator knows when to omit whole parts of a machine in order to reveal concealed mechanisms. The structural engineer strips away all the irrelevant aspects of a building to show the bare bones of a structure. Here we are touching again on the issue of abstraction - the omitting of inessentials and the representation only of what is important for the purpose - which is (or should be) of overriding importance in discussing the application of any technology to Education. It suggests that the effect of VR (and particular kinds of Virtual Environment and Virtual Object) on educational processes and learner perceptions requires more research.

Dynamics in drawing

We have been using the term drawing' for any constructed 3-point perspective representation which is not automated like photography. Turning to a more limited traditional view of drawing' - manual mark-making - there are some special merits that this kind of image has over both photography and computer-rendered models. In many cases, looking at a drawing allows us not just to see the final result, but to decipher the process which produced it.

Productive ambiguities in drawing

We are used to the idea that ambiguity in information is a problem, so the idea of productive ambiguity may be unfamiliar. However, if we consider drawing as a process, rather than a product, it becomes clear that there are merits to imprecision about what is depicted.

Arnheim, in a useful summary (1993) offers some observations on the superiority of sketching over more concretised imagery. These include:

- Sketches are guiding images' whose role as externalisation facilitates the design process rather than being an end in itself
- As long as the guiding image is still developing it remains tentative, generic, vague'
- This vagueness however is by no means a negative quality.' The sketch stands for a whole range of possibilities without being tangibly committed to any one of them.'

The relevance to VR is that, like most computer representations, Virtual Environments and Virtual Objects have a precision and a finality which may be difficult to resist - the real presence' which is one of their strengths could, used inappropriately, be their weakness.

Distortion

Distortion is a loaded word. It implies falsehood. We think we know the difference between a visually distorted and an undistorted image. However, distortion' can be a more subtle affair, which assists rather than impedes the user's experience. As so often, the history of painting provides insights, and Dubery and Willats (1972) give a number of useful examples. We cite two:

- If we survey a wide angle on a scene, using undistorted' 3P-perspective, the resulting image tends to have what we subjectively regard as distortions at the edges - things seem overlarge, or stretched, at the perimeter of the view (Dubery and Willats 1972, p87-89). However, if we distort the objects, as though they were projected onto a curved picture-plane, we create what appears to be a more natural' view (ibid p89-90). In addition, we gain the advantage, as Van Gogh does in his painting Vincent's Room, that viewers sense' their own location in relation to the scene - a key requirement of virtuality. Vincent's Room feels like a small, intimate space, because of, not in spite of, the distortions' imposed on it.

- The work of Cezanne offers another interesting challenge when we think of the relationship between the photographically correct and the feeling real', where accurate spatiality and convincing virtuality seem to make competing demands. Notoriously, Cezanne painted what he saw, in a different sense to that normally meant. Rather than make a painting which is a time-less summation of the whole scene as though apprehended in an instant, Cezanne attended to the process of seeing. Whereas a table-edge which is bisected by a jug is conventionally depicted as the straight line which we know' it to be, Cezanne would rather show the angle that the table-edge seems to make when observed on one side of the jug, and then the angle at the other side, even though, on looking overall, we can see that the two angles are inconsistent - that the lines if extrapolated do not correctly 'join up'.

We are not suggesting that either of these specific techniques should be applied in rendering virtual models to the display, but they are intriguing examples of visual distortion which serve the purpose of making things feel more real.

A more prosaic example is given in Sarkar and Brown (1994), where productive use is made of the distortions of 'fisheye' views. Within the restrictive space offered by a conventional display, their aim was to reconcile the demands of providing maximum contextual information, with the largest, clearest possible view of the current topic of interest. The material to be displayed was a network diagram of linked notes. The solution adopted was to provide a central zone displayed in the standard way, and to progressively compress the image nearer to the perimeter of the screen. Of course, one of the advantages of digitally simulating this 'lens-like' behaviour, is that we can instantly abandon it when it does not serve our purpose, or adopt a view which, for example, compresses only at two edges rather than all four.

The viewpointless view

All the projections we have discussed assume that a viewpoint is necessary, and indeed a view without a viewpoint seems a nonsense. However, children's drawings and some other artforms arguably provide just this: the objects they display carry their canonical views with them as properties of the objects themselves, so that human beings are always displayed full-face, animals are viewed in profile, rainbows are always visible throughout their length. In a limited sense, this is an approach with which computer users are familiar - they see it in an iconic graphical user interface.

Taking liberties with realism

There are a number of attributes of the perceived world which, rather than copying slavishly, we may be able to re-use' for our own purposes, in order to clarify or enhance meaning. Examples include focus, distance, atmosphere and scale.

Focus
We described how the work of Cezanne attempted to capture the process of seeing. Others have used focus in similar ways. It seems certain that Rembrandt's self-portraits owe a part of their extraordinary presence to his depiction of the differential focus of the plane of the face, so that the eyes, the hypnotic subject for any painter observing his own face in a mirror, are in sharp focus and the tip of the nose and the distant parts of the head are relatively defocussed. Baxandall (1985 p80) shows how Chardin uses selectivity in the sharpness and softness of contours, to capture the way in which the eye might prefer certain trajectories across the scene.
In so doing, Chardin almost certainly hopes to lead the viewer's eye on a similar journey in negotiating the parts of his painting. A crude example from this century would be the use of selective focussing to take in and out of attention different parts of a movie scene, perhaps pulling focus' from one character's face to another's.

Distance
Distance is generally perceived as an inconvenience: we have to travel from place to place at a certain speed. Taking advantage of the non-physical nature of virtual reality, we can all but abolish the lapse of time in translating ourselves from place to place. Yet time is an important carrier of information: we use it as meta-information to discern where data (for example sound) is coming from. In place of time-delays arising from the accidents of place, we could use time-delays as an analogue of other qualities, for example the remoteness of one idea (as embodied by an object, building, etc.) from another.

Atmosphere
The effect on clarity and colour of distant objects, is another accident of reality which imposes difficulties and yet which might be turned to advantage. In physical reality, atmosphere is a function of distance and climate: there is nothing to be gained by saddling ourselves with such accidents to no purpose. However we can find uses for the effect. For example, we could assist the user in giving attention to certain features by partially obscuring some and clarifying others. In making such a decision, we are again designing.

The uses of scale
If we took a narrow view of VR as photorealism in 3D, we would deny ourselves all the possibilities of using scale to create meaning. One has only to look at pre-Renaissance painting, or the films of most animators, to see what can be done once unthinking adherence to perspectivally correct' scale is abandoned.

Information spaces: database to Virtual Museum

Virtual environments may have potential as frameworks for knowledge, even where that knowledge is not inherently spatial. Pieces of information may be situated in a virtual space. The analogy with museums has led to the construction of virtual galleries, and certainly non-physical environments do seem to offer some possible benefits:

- Collection of artefacts which used to be housed in one place can be reassembled, perhaps in a reconstruction of their original location. So for example a picture collection dispersed through warfare or natural disaster can be partly reconstituted in a virtual sense.
- Extant collections can be made available in a form which resembles their actual setting to those unable to have access to the real thing.
- Artefacts can be placed in new contexts evoking their original use or setting (something many physical museums struggle to reconcile with scholarship, curatorial care and so forth).
- The emphasis on static, preservable objects associated with most physical museums can be overthrown. Sound and other temporal or fleeting phenomena can be included just as well as physical artefacts.
- New kinds of artefacts can be created to take advantage of the special characteristics of such environments.

However, as always we must beware of any naive assumption that the design of virtual museums is a matter of replicating the style and character of the physical museums with which we are familiar. Three virtual dimensions may be no better than two in representing the structures of knowledge.

Some of the difficulties of virtual galleries are:

- Objects can usefully be considered to occupy a place in as many dimensions as there are metrics for describing them (in database terminology, the space has as many dimensions as there are fields in the record). So for example a painting whose date, author, material, place of origin and main subject are recorded is already situated in a five-dimensional space. Three dimensions are not then enough to do justice to the subtlety of the relationship between each artefact and its neighbours. Of course those dimensions may be reduced to a three-dimensional space for the purposes of display (just as they are in a physical gallery when paintings are exhibited chronologically for example). But this raises another important objection:
- We are faced with a difficult choice. We could place objects in the 3-space we design, and not allow the user to reconfigure the space. In so doing we transfer a major disadvantage of the physical gallery to the virtual one: objects can no more be viewed in new contexts than they could in a physical space. Alternatively, we could allow users to continually re-order the ‘gallery' according to the needs of the moment (for example, on one occasion to see all Italian paintings together, on another, all from the 18th Century, on another all those painted using tempera). Now the risk is that we lose key advantages of virtual environments over conventional databases, particularly the easy memorability of spatial organisation: our cognitive map will be of no use if everything moves around whenever we apply different criteria (Shum 1990). Assuming that virtual environments are more expensive to develop and to deliver, we must ask ourselves what we have gained.

There may also be educational objections to the spatialising of all relations between objects. We have to remember that space in such contexts is a metaphor for other things. Do we risk imposing simplistic structures on the learner's thought by adopting the single dominant metaphor of space? Nelson's remark about metaphor in the interface is relevant here:

> In the metaphorical approach, the metaphor becomes the central concept - the principle, if you will, to which all other aspects of the design must adhere. The problem is [...] that slavish adherence to a metaphor prevents the emergence of things that are genuinely new.
> Nelson, 1990 (original emphasis)

In constructing VRs to house certain kinds of knowledge, we risk tying ourselves to a single model where alternative representations may be better capable of promoting new thoughts. These and related ideas are discussed in greater depth in Gere 1996.

The Homeopathic Fallacy'

Another important issue for educational applications of VR is a difficulty which has been highlighted in relation to hypertext, namely the homeopathic fallacy' (McKendree et al 1995). The objection is to a claim often made: that the graphic visualisation of a relationship between a set of concepts is necessarily educationally productive. Hypertexts map the relation between information objects by one method, and a spatial virtual collection does so by another, but the objection is that in either case the solution is based on a false argument:

1 We intend that a learner should develop a conceptual network of the relations between items of knowledge
2 We depict' such relationships, whether using hypertext links or spatial organisation in a Virtual Environment: essentially the author aims to make the structure manifest on behalf of the learner
3 Now we encounter the problem: the structure is manifest in the system, and we wish it to be in the head of the learner, but we don't possess a magic funnel which allows knowledge to be poured into our heads' (ibid).

A related problem is that of the learner's ability to critically appraise a representation. When we see an N-dimensional information-structure mapped as a series of two-dimensional screens, we are not only able to conceptualise the relations between neighbouring items without those relations being pictured on the display, but also to remember that the sequence as presented is indeed a representation - one amongst many. Could the very persuasiveness of a spatial VR representation be a problem: that learners might cease to see how the relationships between items is contingent, susceptible to revision and questioning - the very sort of realisation which HE should foster in learners?

We must not think that the construction of virtual realities is a common-sense operation in which it is obvious what choices should be made, one in which design has no part. Perhaps there is a danger that we might lose heart in the face of the multiplicity of choices described in this report. But the point is that - to an even greater extent than in the physical world - we can build anything we like. The decisions we make must be informed by a knowledge of the wide range of possibilities open to us, and by an understanding of the processes of perception and cognition. These include not only what we build, but how we represent it and how the user engages with it - above all what educational processes it promotes and why.

Design into production

Preceding sections of this report have deliberately addressed issues which are at several removes from the immediate business of making a Virtual Environment. Here we approach more closely the transition from concept to project.

Some metrics of reality

An unspoken assumption in some discussion of VR is that its technological development is well advanced and that the path before it is unproblematic. We might almost think that VR is within sight of the goal of perfect replication. By contrast, we have emphasised that a VR world is a representation, not a reconstruction of reality. The only representation that would theoretically approach reality would be boundless in size and of infinitely small granularity, modelling the behaviour of sub-atomic particles! This is a practical problem, not a philosophical one, and an important one for education. Not only must those who are building models recognise the limits of the representations they build, but so must the academics and learners making use of them.

Items from a reality checklist
It may be useful to look at a few of the aspects of reality which could in principle be represented using computer models, and notice how only some are possible in VR or even simulation systems.We have tried where possible to describe these metrics in non-computer terms.

Extent and scaling
- Height, depth, breadth

Position and movement
- Location in three-dimensional space; translation, rotation, including non-linear velocities

Passive visual qualities
- Colour: hue, saturation, value; transparency, translucency; reflectivity (matt, lustrous, shiny, gloss, part-mirrored); texture (surface texture, solid texture)

Active visual qualities
- Light source type eg light bulbs, phosphor, lightning; colour of light; distance, direction and spread of light sources

Dynamics and dependencies
- Free movement independent of all other objects; rigid fixing to other objects; able to penetrate some objects, not others; internal object constraints (degrees of freedom); dependent object constraints (within, outside, hinge, slide); hierarchical object constraints (eg. arm moves hand moves finger, and reverse: finger moves hand moves arm)

Physical qualities - passive
- Mass, hardness, brittleness, flexibility, crystallinity eg. splintering on breaking
- Isotropy and directionality
- Insulating/conducting properties for heat, sound, electrical current etc.

Physical qualities - active
- Emitting heat and other forms of radiation

Textural qualities
- Fluffiness, cloudiness

Behaviours
- Tendencies to fall, continue along a trajectory, change colour, decompose, grow, subside, disperse, cool, lose or gain moisture, expand or contract

Animal behaviours and propensities

Human behaviours and propensities
Atmosphere: colour, opacity, humidity, air currents
Responsive sounds
When moved eg. scraping; when struck eg. thud; when squeezed eg. escaping air
Autonomous sounds
eg boiling water, sounds of animate creatures

Even this brief list indicates many aspects of reality which are currently difficult to model in VR systems (together with others which are widely available). Importantly, VR systems have tended - by contrast with various kinds of simulation - have tended to model passive visual qualities, rather than those of activity, behaviour and process.

Before embarking on VR...

Up-to-date technical tutorials can provide advice on the practical steps necessary in order to undertake a VRML project. Here we offer broader questions which should be asked when embarking on a project.

Is the project enhanced by 3D?
Clearly a fundamental issue, this is not always easy to answer. Some help may be provided by the following questions:

- If the project replaces other materials, say student handouts of worksheets, do those existing materials use 3D graphics? If not, why not?
- If the project is one of visualisation, are you proposing more dimensions in the output than in the data? For example it is all too common in business graphics and news media to see two-dimensional data enhanced' by a redundant third dimension - thickened bar charts, pie charts with meaningless heights applied to the segments, and so forth. Similarly, as an aid to navigation a map may do as well as a 3D model, or better.
- Are any graphics needed at all (let alone 3D)? Well designed tables can be as effective as graphics (though the literacy and numeracy of the audience must as always be taken into account).
- Is 3D too much? Perhaps what is proposed is a grand interface to rather trivial content. A number of Virtual Gallery projects have suffered this failing, providing a lavish means of obtaining sparse and finally unrewarding information.

Cost benefit
Can you justify the following requirements in relation to any expected gains?

- Development time: 3D modelling involves significant labour, some of it uninteresting repetitive production work (though this decreases if many similar models are built, since libraries of components can be built up). Currently, output from 3D modellers is not always in perfect VRML form. Debugging of VRML files may be necessary, as may optimisation (for example to reduce polygon count).
- Effort: in addition to perhaps needing to master the relevant software packages for creating 3D models and successfully testing and debugging VRML files, it will be important to maintain and develop the models once built. New versions of browser software may reveal hitherto hidden failings in the model, and user expectation may demand enhancements as new versions of VRML emerge and browsers become more sophisticated. In summary, there may be a need for a commitment of time and effort at

three stages: 1 pre-development, including mastering software packages and techniques; 2 development of models, animations and VRML code; 3 maintenance and enhancement

- Loss of audience: is there a danger that users will be sacrificed who do not have access to suitable bandwidth, hardware and software? For example, developing on-line courses dependent on VRML might be counter-productive, if what is needed is greater access for home-based students.
- Direct costs: in addition to the demands on bandwidth, hardware and software for users, there will be direct costs associated with the machines needed for development.

Perception and cognition

In this report we have commented on the possible dangers of replacing abstract and symbolic forms of knowledge with those which are more direct and concrete.

- Will the loss of abstraction be educationally counter-productive ?
- What will be the effect of 3D for all users? Will all users gain from 3D, or do some have poorer spatial abilities?
- In addition to those cases where 3D is seen as clearly necessary, there may also be other factors. Material may attract attention better than it would do in another form. As long as VR is a novelty, material may be more motivating.
  Spatial metaphors in visualisation, such as mapping time as a physical dimension, have a long history in serving to clarify concepts otherwise more difficult to grasp.
- Points against 3D: It is important to consider whether spatial metaphors (like any other metaphor) will override others models, limiting innovative thought. Will the concretising of concepts work against the benefits of abstraction? Does the user's act of entering the world militate against a critical approach, in which understanding is replaced by wonder?

Designing the world - making the most of 3D

- Are the possibilities of 3D worlds being fully exploited? Are the respective strengths of spatiality and virtuality being fully explored?
- What use can be made of the space as a place for collaboration? What benefits arise from combining the strengths of visualised worlds with those of collaborative exploration, research and projects?

Designing the world - practical constraints

Many of the constraints on VRML models at any given time are contingent on particular versions of software - things which are problematic one month are easily solved the next. However, there are some general guidelines which are likely to remain valid:

- Keep models small: even though bandwidth will increase and computers become more powerful (and optimised for 3D), the need to keep models to manageable size will continue as network traffic also increases, consuming the increasing network capability. There will never be a situation in which models move too fluidly in the browser, or the user's actions are responded to too quickly.

Some simple steps which may be taken include:

- Use in-line instances where possible: to avoid unnecessary downloading, build in a modular way where possible, reusing predefined components, so that each is downloaded only once and used repeatedly.
- Avoid texture maps if possible: texture maps are data-hungry compared with the geometry of objects, and the computation of the mapping slows the rendering process.

- Test the use of LODs (levels of detail) with various browsers. With LODs, navigation may be stop-go, travelling smoothly while a model is in RAM and then lurching when the next LOD is loaded. However, without LODs, navigation may be slow all the time.
- Though an increasing number of modellers can produce VRML, test models and effects (such as transparency) at an early stage, to avoid wasted effort. Put all models through economisers and parsers to reduce their size and check their validity.

Designing the world - managing the process

The design of a VR world can be, and should be, almost as difficult as real architectural design. Indeed, as modelling and modellers become more sophisticated, their construction will increasingly resemble architecture (though not necessarily visually). For example, as third-party libraries become increasingly available, there will be problems with consistent quality and maintenance. Likewise management issues will become important as model-building is farmed out to large teams, some of them sub-contracted. And issues of control and negotiation will arise as worlds are connected to each other and virtual worlds come to be seen as public spaces.

VRML now

VRML is the standard being developed for putting virtual realities onto the Internet. For many in higher education it will be with this set of technologies, rather than any other, that they will experience virtual reality.

VRML, along with the Internet in general, offers a number of attractive features for the higher education sector:

- It has cross-platform compatibility
- Much of the software described below can be downloaded for free
- As VRML sits upon existing World Wide Web tools, existing student knowledge of these is applicable, easing use.

The costs of fully immersive VR equipment, with head-mounted displays and specialised input devices, is likely to make them prohibitive for all but the most specific applications. An interesting experimental result, pertinent to this point, comes from Satalich (1996) who looked at the claim that working within a virtual environment would teach the user to navigate naturally in that space. In fact, it was found that a group who only studied a map of the space performed better than those who had 30 minutes of prior exposure to the same space.

Satalich suggested that there may be a number of reasons for these results, such as the novelty of the technology for its experimental subjects, or the limited time during which they where immersed. These concerns are certainly worth bearing in mind, but it does suggest that until they are resolved, we need not look on desktop systems as necessarily inferior. This is especially important when one considers the expense, technical difficulties and health concerns associated with head-mounted displays (Jones 1996).

What follows will concentrate on VRML-based desktop VR, but even this section must begin with a warning: VRML is new and is in a state of rapid development. Details outlined below will change, and URLs given may not remain valid for long. However, the major Web sites, such as the VRML Repository, will almost certainly still be present to give the most up-to-date information available.

This situation of constant change is a problem, as the technologies often do not work fully, but is also very exciting, as ideas that are proposed one month appear as beta versions of software soon after. This speed of development has the consequence that, although the version 2 of the VRML standard is not yet fully defined as of the time of writing, we already need to look more at its potential rather than at what is currently possible with version 1. Doubtless, by the time many have read this report, fully functional VRML-2 tools and worlds will be available. There is, of course, some risk with this approach, but our experience of the VRML community in the last year gives us confidence. The take up and interest in the subject has grown enormously, and has been matched by the development of tools. The rapid pace of change makes the Internet - rather than one-off reports - the best source for information on the subject.

It is worth a brief digression to understand why VRML has suddenly taken off in the past year. This is no doubt partly due to the perceived inherent excitement of virtual worlds. One only needs to look at the success of the cyberpunk novels Neuromancer and Snow Crash. Although the reality of the technology is very far from the fictions, the excitement generated by them is real. Much of what is fiction will remain so, but many ideas can also be expected to come to fruition in the near future following the efforts of major players in the computer industry such as Netscape, Sony, IBM, Intel and Microsoft, now actively interested in 3D across the Web.

The importance of the term virtual reality' must be seen in a different context for VRML than is the case for much of the advanced, high cost, virtual reality systems. The aim of an Internet-based technology is to make the three-dimensional worlds and objects as accessible to as wide a public as possible. The models created must be able to be seen on a range of machines that are readily available to the average user. This is particularly true of many HE institutions which often have many less powerful machines in situ. Speeds of network transmission for models also pose another problem. So for the immediate future it is likely that we must satisfy ourselves with simple (that is, small) models. They are unlikely to approach what anyone might call 'reality' for a long time. The models must, therefore, serve other purposes to be effective and useful. But as we shall see, they can convey abstract qualities, and in the near future can provide 'places' for people to meet, and these are sufficient reasons to engage in the technology.

Information on the current VRML standard are covered by another AGOCG:SIMA report, which can be obtained online (Ashdown 1996). We shall therefore concentrate on some of the issues and concerns that have arisen from our own use of VRML, and where we see its potential.

VRML version 1
The aim of VRML is to bring to the Internet the advantages of 3-D spaces, known in VRML as worlds whether they comprise environments or single objects (and using the file suffix .wrl). These are built to be shared between widely distributed users. Just as the standard WWW browser allows the user to download and display pages containing texts and images that have been appropriately marked-up, so a VRML-compliant browser allows for the display of 3-D models that have been described in VRML. Using the mouse and/or keys, the user can move their viewpoint through the space.

There are two main classes of browser: stand-alone, or a helper application associated with a standard Web browser such as Netscape Navigator or Mosaic. At the date of writing , in our own experience, there are various issues and problems related to browsers.

Many browsers are not fully compliant and fail to display models which contain advanced features. This situation has probably arisen due to the perception that VRML-1 is going to be a temporary measure. Its functionality is limited, and seems more a proof of concept. To overcome limitations various browsers, such as Live3D, support additional features. Equally, many browser authors now unsurprisingly seem more concerned with working on version 2-compliant browsers, than with continuing to improve those for version 1.

Secondly, there is a problem with browser authors porting software to the Macintosh. Most of the browsers run on IBM compatible machines, with some for high-end machines (such as Suns and SGIs), but the Macintosh is not well supported at present. However many university faculties, especially those in art and design, have a heavy investment in Macintosh equipment, and it can only be hoped that this situation will change. The appearance of Live3D from Netscape on the Macintosh has slightly improved matters, but the basic problem remains.

VRML-1 is simple, in that it only allows for the creation of static worlds. The interaction with these worlds consists of navigating through this space, and the equivalent of WWW hotspots on objects within the world, which allow the user to move to other worlds - or to any other Internet feature such as web pages, ftp sites etc.

VRML is based on a subset of Open Inventor from Silicon Graphics, and it defines a set of objects and functions for modelling simple 3D graphics. These are known as nodes, which are arranged in hierarchies called scene graphs. There is a top-down arrangement in which nodes that are described earlier in a scene affect later ones, but this can be limited by the use of separator nodes. This allows one to, for example, colour one part of the scene, but prevent the colouring affecting following nodes. A VRML file is an ascii file which is interpreted by the browser and converted into a 3-D display of the described world. For example, the VRML author can specify the geometric shape of objects in the world, transformations on them, surface colouring and textures, how the world is lit, and the cameras which create the final display for the user (see Ashdown 1996 for details).

Due to the need for the files to pass through the Internet, and the need for the files to be displayed on low-end machines, there are size issues for VRML which have been addressed in the language design. Through in-lining, large VRML scenes can be broken up into smaller files which are then loaded as required.

Another feature - one that has been used by games developers for many years - is levels of detail, or LODs, which assist slower machines. If we are far away from an object, the browser can display a simplified version of it. As we move in closer this can be dynamically updated to display more detailed versions of the same object.

Authoring VRML 1.0 worlds
There are various ways in which the VRML files can be created:

- Worlds can be created from scratch with a text editor, typing in the VRML language directly and then loading the final file into a browser for viewing. This is not practical if the world design is very complex and contains non-primitive shape forms.

- There are some simple tools specifically for designing VRML worlds, such as Caligari's Pioneer (http://www.caligari.com/com/products/pfeat.html). This has specific support for VRML-1 features such as inlines, LODs and links.

- There are application packages which allow the author to design the world on-screen graphically and which then output a VRML-compliant file. For example, both 3D-Studio and Strata Pro Blitz have options to save files in the VRML format. It is also certain that other popular graphical tools will offer VRML output, as they come to appreciate the emerging importance of the standard.

- There are a number of converter programs which will take file formats created by other graphics packages, and convert then to VRML. These packages are particularly useful depending on the type of object or world being created. That is, each discipline - architecture, engineering or chemistry, for example - has its particular needs and methods of working which are supported by its own preferred modelling packages. Often these features are not to be found in any one tool, but through translators the model can be made available in VRML format. Many of these converters/translators are available free online, and can be found via the VRML Repository (see Internet Resources, page 97).

- It is also possible to have worlds generated on the fly' by a program running on the server. This is likely to be an area of growing importance as VRML is used, for example, to display data from a range of analytical programs. Postgraduate students at the Centre for Electronic Arts are currently investigating the construction of an organic' element for a world, which will grow according to the number of Web hits' that the world receives.

Despite the obvious value of GUI tools, some knowledge of the VRML language has proved valuable for many current users, who have found it useful to build a few simple worlds using this method as a means to understand the general principles of the standard. It is also useful where minor errors are detected in the file output by a GUI tool, since these can be simply edited in a text-editor. Similarly, application packages often do not fully implement the standard, since they are generally a modification of an existing tool. For example, linking objects to URLs may not be supported. Therefore the ability to get into the file created by an application to edit it as required may still be necessary in the short term.

The value of this, however, is likely to diminish as both browsers and graphics packages begin to fully implement the standard.

The Future of VRML

The current VRML standard is one that concentrates on static objects and scenes which users can move around. As this report is being written the debate and process for the specification of VRML-2 is being undertaken, managed by the VRML Architecture Group (VAG). A number of mainly commercial groups are developing software for interactive social worlds as a test bed, elements of which they hope will contribute to the future standard. The future is in two areas, which are increasingly related.

Dynamic Models

VRML-2, also known as Moving Worlds, provides, as its name suggests, for the introduction of action in the modelled world. Either objects in the world can act and react to each other under program control, or they can respond to the user's actions in some way. For educational purposes, this will in many cases be essential. Byrne (1996) in a controlled experiment found that immersive VR seemed to offer little advantage over a Window-on-the-World (WoW) approach, but interactivity perhaps did. This was specifically in an example where the technology was used to support education. She sought many reasons to explain the data, but the answer may be that immersion may not be required to achieve most of the benefits to education. The results seemed to suggest that it was the interactivity that was important. This corresponds to a constructivist theory of learning, in which the student builds and manipulates problem spaces, rather than being merely an observer.

Social Worlds

One of the most important aspects of the Internet is its ability to facilitate interpersonal communication. IRC and MOOs are examples of such facilities that are very popular. These may prove exemplary for the development of applications for VRML. The history of the Internet is largely one of the victory of chatter over content, but like the telephone the social worlds will be very important. In the HE sector, it will remain the case that much learning is achieved through student-student and staff-student communication: social worlds can facilitate this.

We will look at these two developments, Moving Worlds and Social Spaces, in turn.

Moving Worlds

Moving Worlds extends VRML in various important ways. The VAG (http://vag.vrml.org) proposes (May 1996) that the Moving Worlds specification goes through a series of drafts with a release date of August 4th, timed to coincide with SIGGRAPH 96. However, Silicon Graphics (http://vrml.sgi.com) has already released a beta version of a browser, Cosmo, that implements many of the features of Moving Worlds (VRML 2.0). Clearly, from this response, it is almost certain that a fully compliant browser will be available very soon after the standard has been fixed. For many readers, therefore, it is probably worth jumping over VRML version 1, straight to version 2, the gain in functionality outweighing the risks of working to forthcoming standards.

The features that Moving Worlds currently includes, most of which are likely to remain, are:

- International character sets for text can be displayed using UTF-8 encoding.
- A set of new nodes has been added to increase the realism' in models that are intended to represent the outdoor world around. It will be possible to create ground and sky backdrops, adding distant mountains and clouds, for instance. The effect of distance can be further enhanced with fog effects. In addition, there will also be the option of defining an irregular terrain rather than being constrained to a flat ground plane.
- Sound generating nodes will also enhance the sense of realism. Thus, doors can creak, phones ring and dogs bark. Intel has announced “Realistic Sound Experience” (http://www.intel.com/ial/rsx) which will allow spatialised sound to be put into VRML worlds, such as the DirectedSound and pointSound nodes in Moving Worlds.
- New sensor nodes will set off certain events when one enters specific areas, or click on certain objects. So, for example, as you approach an object it can be triggered to start some action or make a noise.
- Collision detection ensures that objects can act as if solid. That is, the user, or other object will not go through walls and floors.
- Script nodes allow for the animation of objects in the world.
- Multi-user environments. There are many approaches to creating multi-user worlds, and Moving Worlds aims to provide the functionality required for these, but without dictating which approach is to be used.

Moving Worlds is very extensible, and fully-integrated Java applets can be used to create motion and enable interactivity: there is a small tutorial on what can be achieved at http://www.javaworld.com/javaworld/jw-03-1996/jw-03-javavrml.html. JavaScript allows scripted communication between objects both within and without the virtual world. It is this aspect of the Moving Worlds definition, however, which is still uncertain at the time of writing. Furthermore, Plug-ins within the 3D worlds will facilitate such services as streaming audio and video, and slide presentation.

Social Worlds using VRML

There is a series of worlds being made available, experimenting with increasing functionality, mostly based on the VRML format. They include either scripting and construction facilities, or additional communication, be it textual, audio or even video. Most of these take VRML models and allow them to be used within multi-user shared environments. That is, users from anywhere in the world can connect to these, inhabiting them embodied in an avatar (a 3D character, also usually modelled in VRML). Users can then meet' other users and chat. In VRML-2 the users can also manipulate the world.

A list of such shared worlds, or habitats as they are often known, is currently kept by Electric Communities at http://www.communities.com/habitat.html. For the current discussion, some of the most important ones are:

- CyberGate

Produced by Black Sun Interactive (http://www2.blacksun.com/about/index. html). This is one of the most popular of the new shared worlds. You can create a VRML model on your own server, which can then be made accessible to others who have a CyberGate browser, and become a shared space.

- OnLive

Unlike most others this allows for voice communication between participants (http://www.onlive.com/utopiamkt/index.htm).

These projects indicate the kind of features that are now coming available, some of which will make the next version of the standard. In addition to multi-user access, we have interaction and animation. The ability for the world to carry both audio and video is also shown in some, such as InterSpace by NTT.

The importance of these worlds is indicated by the move of some of the largest players in the computer field into the area, such as:
Microsoft with V-Chat (http://www.msn.com/v-chat/index.htm)
Sony's CyberPassage (http://vs.sony.co.jp/VS-E/works/server/swelcome.html)
IBM's Virtual World (http://www.software.ibm.com/software/virworld.html) and
Intel's Moondo (http://www.intel.com/iaweb/moondo/index.htm).

There are some popular habitats, such as Worlds Chat and AlphaWorld created by Worlds Inc., that are not true VRML. But they too have recently announced that their next technology, Gamma, will be VRML-2 compliant.

It is our belief that these shared worlds, based on VRML, will be very important in the future. One lesson from history is that the communication aspects of any technology become important and are often overlooked in the early days (the telephone was originally seen as a good way of listening to opera without visiting the theatre). The ability for students to not only visit another VRML to explore it, but also to possibly meet' its authors is very exciting.

Many current uses of VRML-1 are simple at best, due to the constraints described above. In many cases they are existing models created for other reasons, and simply converted to VRML. Proper applications that are designed to fully use VRML are still in short supply, and there is a definite sense that everyone is awaiting VRML-2 to be able to create models that have some real value. Pointers to some of these existing applications can be found at the VRML Repository.

Future roles of VRML in Higher Education

The possible applications to higher education can be divided into three areas: as a medium to enhance teaching and research within the HE sector, as a new subject of study, and as a collaborative space. We concentrate, however, on the latter two as they are less often discussed in the context of HE, where we believe they will prove important.

1 As a medium for education & research

The Internet generally is being proposed as a means to facilitate both teaching and research in the tertiary sector. Indeed this is the most obvious role for VRML. Three-D content can be created and distributed to one's own students, or made more widely available to learners across the globe - it functions as a medium.

In its simple, extant version 1, VRML is essentially a means of graphic presentation. Therefore, any textbook or course materials that have 3-D graphics, such as architecture, chemistry and engineering, may benefit from VRML. For example, the publishing of chemical models can be assisted as suggested by Casher & Rzepa (1996). If the course materials, or the research papers, require multiple drawing of a 3D object or space, then clearly VRML will give benefits.

More abstract representation - that is, data visualisation - is also increasingly supported. VRML modules are available in visualisation systems such as Data Explorer from IBM, and IRIS Explorer for NAG. Data visualisation can be easily supported in that the objects can use the WWWAnchor node to point to textual, tabular material describing the data in other forms. This ability will extend the range of subjects that can benefit from VRML to include those that deal not with spatial objects, but call upon quantitative or qualitative data that can be spatially represented, such as the social sciences

Rather than list them again here, we direct the reader to the VRML Repository list of example applications.

2 As a new subject

New technologies, such as multi-media and computer graphics have provided the basis for new professions and courses. Equally, it is likely that a fully functioning VRML technology will provide opportunities for courses in the design and implementation of virtual spaces, be they for social use or data visualisation purposes. There will be a requirement for specifically trained designers who understand the issues associated with virtual spaces.

It is interesting to consider what elements will be required for such a course. There is, naturally, a requirement for an understanding of the underlying technologies: 3D modelling, the Internet and their associated tools. But there is also a need for specific design skills. It is to discussion of this broader kind of understanding that this report has been largely given over.

Many of the skills of thinking in terms of space, and human relationships to space, are to be found in architecture, and there is some evidence that many architects in the Unites States in particular are moving into the design of virtual spaces (Phillips Mahoney 1995). In another example, the effective use of coloured lights on coloured surfaces within worlds is one that comes less from traditional graphic design and more from theatre set design or interior design. To understand collaborative, multi-user spaces suggest that some psychological and sociological content will also be needed.

The employment market for the graduates of such courses is manifold. There is a games market developing which also requires trained 3D world designers. For example, Damage (http://www.dircon.co.uk/balanda/damage. html) allows players to become part of the cast in a developing game environment. Based on a client/server Internet model, players from all over the world play together either in opposition or cooperatively. With programmable behaviours, and a heavily modifiable character toolkit, players can change or barter their attributes to let them cope with the different scenarios that the game throws up. They intend to be compatible with the VRML-2 standard.

Another such game is Westworld 2000, by Byron Preiss Multimedia (http://www.byronpreiss.com/brook/westwld/west.htm) - a multimedia CD-ROM game with on-line multi-user play. The game lets the user explore the three-dimensional virtual theme parks, while interacting with other players and three-dimensional humanoid robots. The characters and environment respond dynamically to the player's actions.

In addition, it is likely that many organisation will create permanent, shared, online worlds to support, for example, customer enquiries, or collaborative working. These too will require specialist designers. Recently Cybergate saw its first political meeting, where a space was created to allow a US Congressional candidate to meet' the public, indicating range of applications that are being tried.

3 As a collaborative medium

As already described above, one of the longer-term, but possibly most exciting, opportunities lies in the educational use of the shared spaces for distance teaching and project collaboration. The communicative worlds similar to those described above, would allow for students and staff to engage in distributed meetings, seminars and tutorials. A student at the Centre for Electronic Arts is currently experimenting with a CyberGate-compliant space that is used for tutorials (http://kiad.ac.uk/vrml/zone1.wrl)

A number of trials using the existing technologies are being tried. For example, the Virtual Online University is experimenting with Chaco's Pueblo browser which includes VRML, as an interface for distributed learning (http://www.chaco.com/press/vou1.html).

Another experiment which did not use VRML, but gives an idea of where a future VRML standard is heading, is the CyberCampus. This system was demonstrated by NTT and employs Sense8 virtual reality tools (http://www.sense8.com), which provides a multi-user space. The College of Marin created an educational program in which students along with a teacher walked' around a model of an actual geological location, which had pointers to teaching materials.

Increasingly the Universities seek to adapt, supporting learning-for-life', in addition to more conventional students. Such technologies may mitigate some of the problems of isolation that distance learning brings. The students not only access course materials, but can also meet in social spaces to discuss them with staff and fellow students.

Recommendations

1 A variety of courses in Higher Education should deal with VR as a rich and complex subject, offering a wide range of conceptual and technical approaches.

Disciplines and departments suitable to make a contribution include (but are not limited to) Computer Science, Art and Design, Psychology, Sociology, Critical Studies.

In future, VR is likely to offer new opportunities for study and creative work to disciplines as diverse as Ethnography or Drama, in the sense that VR can be the subject and the site of these studies, not simply the medium by which teaching and learning may take place.

Single-discipline courses dealing with VR should aim to inform themselves and their students of the possible contribution of other disciplines, through reading-lists (including URLs) and the provision of AGOCG SIMA materials.

It is not too soon to establish Masters programmes specifically geared to the design of networked spatial objects and environments.

2 VR projects should aim where possible to be strongly multi-disciplinary. The construction of Virtual Environments should be seen as a design problem in the broadest sense. Teamwork will be necessary (and this is a skill which should be built into the teaching of the designers).

3 VR project teams should take advantage of the fact that there are few rules' as to what a Virtual Environment should be like. But they should also aim to harness the wide range of knowledge embodied in precursor and contributory disciplines such as Fine Art, Literature, Perception Studies, Architecture, Product Design, Theatrical Design and others.

4 VR projects with an educational remit should recall at all times the criteria of educational effectiveness. The drawbacks as well as the advantages of various aspects should be evaluated, including the effects of concretising information, of appropriateness of spatial metaphors and of the interface.

5 Regarding shared worlds, they should be investigated to support distance tutoring and research, but with caution - to ensure that the social grouping effect of actual meeting is not lost. They should supplement real tutorial and seminars, not replace them (at least at present until the issues are resolved).

6 A multiplicity of models and worlds is required, in order for a threshold of spatial literacy' to be reached. The making of models by different disciplines for different purposes should be encouraged, and the results put into the public domain through the medium of VRML.

7 Creators of VR from all disciplines should seek to influence new standards as they emerge.

8 Additional wide-ranging research should be undertaken into many aspects of VR, especially looking at both narrow and broad effects on informational and educational effectiveness.

9 Virtual reality and reality - five guidelines for designers:

a Do not mimic reality without good reason.

b Be selective - use detail only where it serves a purpose.

c If reality is required, suggest it by carefully chosen items together with simple forms.

d Harness those aspects of reality which serve the purposes of the project, suppressing those which do not.

e Feel free to mix realism with non-realism.

References

Addams R 1834 'An Account of a Peculiar Optical Phenomenon Seen after Having Looked at a Moving Body' The London and Edinburgh Philosophical Magazine and Journal of Science (5) pp373-374

Anderson RA 1989 'Visual and eye movement functions of the posterior parietal cortex' in Cowan WM, Shooter EM, Stevens CF and Thompson RF (eds) Annual Review of Neurosciences (12) pp377-403

Ames A L, Nadeau DR & Moreland J L 1995 The VRML Sourcebook John Wiley and Sons

Arnheim, Rudolf 1993 'Sketching and the psychology of design' in Design Issues Vol IX No 2 Spring 1993

Ashdown, N. 1996 The Virtual Reality Modelling Language in Art and Design Higher Education http://info.mcc.ac.uk/CGU/SIMA/vrart/title.html

Bablet, D 1962 Edward Gordon Craig Published in French 1962, Translation Heinemann 1966

Bachelard, Gaston 1958 La Poetique de l'Espace, translated as The Poetics of Space by Orion 1964, republished Beacon Press, Massachusetts 1994

Bann, Stephen 1987 'Art' in Cohn-Sherbrook, Dan and Irwin, Michael (eds) Exploring reality Allen and Unwin p83

Barell, J 1982 'Geographies of Hardy's Wessex' in Journal of Historical Geography 8 p347-361

Baxandall, Michael 1985 Patterns of intention - on the historical explanation of pictures Yale University

Berger, John 1972 Ways of Seeing BBC/Penguin London

Berthenthal BI, Profitt DR, Spetner NB and Thomas MA 1985 'The development of infant sensitivity to biomechancial motions' Child Development (56) p531-543

Bolter, J David 1984 Turing's man - western culture in the computer age North Carolina University Press (also Penguin 1986)

Boring EG (1942) Sensation and Perception in the History of Experimental Psychology Appleton-Century-Crofts, New York

Brown, C Marlin 1985 Human Computer Interface Design Guidelines

Bulmer, Martin GS 1995 The application of the psychological theory of flow to user interface design Unpublished essay, June 1995, MA Design for Interactive Media, Middlesex University, UK

Burr and Ross 1986 Visual prccessing of motion' Trends in Neurosciences (9) 7 pp304-307

Byrne, CM 1996 Water on tap: The use of Virtual Reality as an educational tool PhD Thesis, Human Interface Technology Lab, University of Washington http://www.hitl.washington.edu/publications/dissertations/Byrne

Caelli T 1981 Visual perception: theory and pratice Pergamon, London

Campbell PF 1979 ‘Artistic motion cues, number of pictures, and first-grade children's interpretation of mathematics textbook pictures' Journal for Research in Mathematics Education (10) 2 pp148-151

Casher, O & Rzepa HS 1996 The Molecular Object Toolkit: A new generation of VRML visualisation tools for use in electronic journals Proc. Eurographics-96, Jones H & Vicars D (eds) London, 26-28th March pp173-178

Csikszentmihalyi, M 1992 Flow: the psychology of happiness Rider, London, UK

Csikszentmihalyi, M and Csikszentmihalyi IS 1988 Optimal Experience: Psychological Studles of Flow in Consciousness Cambridge University Press, UK

Csikszentmihalyi, M. 1975 Beyond Boredom and Anxiety Jossey-Bass Inc. San Francisco, USA

Corbetta M, Miezen FM, Dobmeyer S, Shulman GL et al 1990 ‘Attentional modulation of neural processing of shape, colour and velocity in humans' Science (248) 4962 pp1556-1559

Cutting JE, Proffitt DR, Kozlowski LT and Johansson G 1977 ‘Studies on visual perception of locomotion' Perception (6) 4 pp365-376

De Bruyn B and Orban GA 1990 ‘The role of direction information in the perception of geometric optic flow components' Perception and Psychophysics (47) 5 pp433-438

DIMPE 1990 DIMPE Project - European multi-media publishing in the 1990s (8 pages) Mirror Group Newspapers, London

Dittrich WH 1993 ‘Action categories and the perception of biological motion' Perception (22) 1 pp15-22

Dodwell PC 1983 ‘The Lie transformation group approach to visual perception' Perception and Psychophysics (34) pp1-16

Dubery, Fred and Willats, John 1972 Perspective and other drawing systems The Herbert Press, London

Dupont P, Orban GA, DeBruyn B and Verbruggen A et al 1994 ‘Many areas in the human brain respond to visual motion' Journal of Neurophysiology (72) 3 pp1420-1424

Edmonds, Ernest A 1992 The separable user interface Academic Press, London

Felix N 1995 Analysis of spatial representation from traditional techniques to computer graphics Unpublished PhD thesis, Strathclyde University, Glasgow

Fitzpatrick G, Tolone W J & Kaplan SM 1995 ‘Work, Locales and Distributed Social Worlds' in Proceedings of the Fourth European Conference on Computer-Supported Cooperative Work.

Marmolin H, Sundblad Y & Schmidt K (eds), Kluwer Academic Publishers, pp. 1-16.

Foucault, Michel 1974 The order of things (originally published as Les mots et les choses, Gallimard) Routledge, London

Friedman SL and Stevenson MB 1975 ‘Developmental changes in the understanding of implied motion in two-dimensional pictures' Child Development (46) pp773-778

Gere, Charles 1996 The Virtual Museum PhD Dissertation, Middlesex University, UK

Gibson, JJ 1979 The ecological approach to visual perception Houghton Mifflin, New York

Giddens, A. 1984 The Constitution of Society. Berkeley: University of California Press

Girouard, Mark 1978 Life in the English country house Yale University Press

Gombrich, EH 1979 The sense of order: a study in the psychology of decorative art Phaidon, Oxford

Hall, Edward T 1959 The Silent Language Doubleday, USA

Hall, Edward T 1966 The Hidden Dimension: man's use of space in public and private The Bodley Head, London (1969)

Harris M, Freeman T and Williams G 1992 ‘Surface layout and retinal flow' in Humphreys GW (ed) Understanding vision Blackwell, Oxford pp199-231

Harris MG 1994 ‘Optic and retinal flow', In Smith AT and Snowden RJ (eds) Visual Detection of Motion Academic, London pp307-332

Hildreth EC and Koch C 1987 ‘The analysis of visual motion: from computational theory to neuronal mechanisms' Annual review of neuroscience (10) pp477-533

Hillier, B and Hanson, J 1984 The social logic of space Cambridge University Press

Hoffmann, WC 1970 ‘Higher visual perception as prolongation of the basic Lie transformation group' Mathematical Biosciences (6) pp437-471

Hoffmann, WC and Dodwell PC 1985 ‘Geometric psychology generates the visual Gestalt' Canadian Journal of Psychology (39) 4 pp491-528

Hopgood, FRA 1993 ‘Use of time and space in multimedia systems' in Multimedia in Higher Education: portability and networking AGOCG Technical Report No. 24 December 1993

Hunter, William 1932 Scrutiny of cinema Wishart & Co, London

Hutchins, Edwin L, Hollan, James D and Norman, Don A 1986 ‘Direct manipulation interfaces' in Norman, Don A and Draper, Stephen W (eds) User centred system design: new perspectives on human-computer interaction Lawrence Erlbaum Associates, Hillsdale, NJ pp87-124

Jackson, Peter 1989 Maps of meaning Unwin Hyman 1989, reprinted Routledge, London and

New York 1992

Johansson G 1973a ‘Visual perception of biological motion and a model for its analysis', Perception and Psychophysics (14) pp201-211

Johansson G 1975 ‘Visual motion perception', Scientific American (232) pp76-88

Johansson G 1995 ‘Personal Comments' in Johansson G, Bergstrom SS and Epstein W (eds) Perceiving events and objects: resources for ecological psychology Lawrence Erlbaum Associates, Hillsdale pp495-503

Johansson G and BÜrjesson E 1989 ‘Toward a new theory of vision: studies in wide-angle space perception' Ecological Psychology (1) pp301-331

Jones DH 1996 ‘Virtual reality: panacaea or Pandora's box?' Virtual Reality (3) in press

Juola JF and Breitmeyer BG 1989 ‘A discussion of models of motion perception', In Elsendoorn BAG and Bouma H Working Models of Human Perception Academic Press, London pp251-259

Kanizsa G and Gerbino W 1976 ‘Convexity and symmetry in figure-ground organisation', in Henle M, Vision and Artifact, Springer, New York

Katz, Ephraim 1994 The Macmillan International Film Encyclopedia Pan Macmillan, London

Kaufmann F 1995 ‘Development of motion perception in early infancy' European Journal of Pediatrics (154) 9 pp548-553

Kellman PJ 1995 ‘Ontogenesis of space and motion perception', in Epstein W and Rogers S (eds), Perception of Space and Motion (2nd Edition), Academic, San Diego pp327-364

Kellman PJ and Spelke ES 1983 ‘Perception of partly occluded objects in infancy', Cognitive Psychology (15) 4 pp483-524

Koenderink JJ and Van Dorn AJ 1976 ‘Local structure of movement parallax of the plane' Journal of the Optical Society of America (66) pp717-723

Kozlowski LT and Cutting JE 1977 ‘Recognising the sex of a walker from a dynamic point light display' Perception and Psychophysics (21) 6 pp575-580

Krippendorf, K 1989 ‘On the essential contexts of artefacts or on the proposition that “Design is making sense (of things)”' in Design Issues Vol V No 2 Spring 1989

Krumhansl CL 1984 ‘Independent processing of visual form and motion' Perception (12) 4 pp417-423

Lansdown, John and Schofield, Simon 1995 ‘Expressive rendering: a review of nonphotorealistic techniques' in IEEE Computer Graphics and Applications

Laurel, Brenda 1991 Computers as Theatre Addison-Wesley

Laurillard, Diana 1993 Rethinking University Teaching - a framework for the effective use of educational technology Routledge

Macdonald, Gus 1979 Camera - a Victorian eyewitness Batsford, London

Markus, Thomas A 1993 Buildings and power - freedom and control in the origin of modern building types Routledge, London

Mather G and West S 1993 ‘Evidence for second order motion detectors' Vision Research (33) 8 pp1109-1112

McKendree, Jean; Reader, Will; Hammond, Nick 1995 ‘The homeopathic fallacy' in learning from hypertext Interactions, July 1995, ACM SIGCHI, New York

Miles, John 1987 Design for desktop publishing: a guide to layout and typography on the personal computer John Taylor Book Ventures, Hatfield UK

Nakayama K 1985 ‘Biological image motion processing: a review' Vision research (252) 5 pp625-660

Nelson, Theodor H 1990 ‘The right way to think about software design' in Laurel, Brenda (Ed) The art of human-computer interface design Addison-Wesley, pp239

Norman, Don 1988 The psychology of everyday things Basic Books, New York

Ong, Walter J 1982 Orality and literacy - the technologizing of the word Methuen (reprinted Routledge 1988, London and New York)

Passini R 1992 Wayfinding: people, signs and architecture McGraw-Hill Ryerson

Palmer SE 1992 ‘Modern theories of Gestalt perception', In Humphreys GW, Understanding Vision, Blackwell, Oxford pp39-70

Pavlova MA 1992 ‘Preschoolers can recognise biological motion patterns', Perception: Supplement 2 on European Conference on Visual Perception (21) p20 (Abstract only)

Pesce, M. 1995 VRML - Browsing and Building Cyberspace New Riders Publishing

Pevsner, Niklaus 1943 An outline of European architecture Pelican

Phillips Mahoney, D. 1995 ‘Architecture for the Fun of It' Computer Graphics World December 1995, pp36-45.

Pinker, Steven 1994 The Language Instinct Penguin

Postman, Neil 1982 The disappearance of childhood Delacorte Press, New York

Postman, Neil 1985 Amusing ourselves to death - public discourse in the age of show business Originally published USA; published in Great Britain by Heinemann 1986

Runeson S and Frykholm G 1983 ‘Kinematic specifications of dynamics as an informational basis for person-and-action perception: expectation, gender-recognition and deceptive intention'

Journal of Experimental Psychology: General (112) p585-615

Sarkar, Manojit and Brown, Marc H 1994 ‘Graphical Fisheye Views' in Communications of the ACM Vol. 37 No.12 December 1994

Satalich GA 1995 Navigating and wayfinding in Virtual Reality: finding proper tools and cues to enhance navigation awareness Masters Thesis, Human Interface Technology Lab, University of Washington
http://www.hitl.washington.edu/publications/satalich/

Schenk T and Zihl J 1995 ‘Motion perception after brain damage' Perception (24) Supplement p54 (Abstract only)

Sekuler R 1975 ‘Visual motion perception', in Carterette C and Freeman MP (eds), Handbook of Perception (Vol 5), Academic, New York pp385-430

Semper, G 1989 The four elements of architecture and other writings (translated Mallgrave HF and Herrmann W) Cambridge University Press

Serres, Michael 1977 ‘Language and space - from Oedipus to Zola' in Harai, Josue V and Bell, David F (eds) Hermes: Literature, Science and Philosophy John Hopkins University Press, Baltimore 1983

Shneiderman 1992 Designing the user interface - strategies for effective human-computer interaction 2nd Edition Addison Wesley

Shum, S 1990 ‘Real and virtual spaces - mapping from spatial cognition to hypertext' Hypermedia (2) 2 pp133-158

Smith AT and Snowden RJ (1994) Motion ‘detection: An overview', In Smith AT and Snowden RJ (eds) Visual Detection of Motion Academic, London pp3-13

Sontag, Susan 1977 On Photography Penguin

Sumi, S 1984 ‘Upside down presentation of the Johanssson moving light spot pattern' Perception (13) 3 pp283-286

Thimbleby, Harold 1991 ‘Can anyone work the video?' New Scientist Vol 129 1757 pp48-51

Tufte, Edward R 1990 Envisioning Information Graphics Press, Cheshire, Connecticut

Ullman S 1987 ‘Visual motion perception', in Adelman G (ed) Encyclopedia of Neuroscience Birkhèuser, Boston pp1280-1281

Warren WH and Hannon DJ 1988 ‘Direction of self-motion is perceived from optical flow' Nature (336) 6195 pp162-163

Warren WH 1995 ‘Self-motion: visual perception and visual control' in Epstein W and Rogers S (eds) Perception of space and motion (2nd Edition) Academic Press, San Diego pp263-325

Whorf, BL 1956 Language, thought and reality MIT Press

Williams, Rosalind 1990 Notes on the Underground: an essay on technology, society and the imagination MIT Press

Yates, Frances 1966 The Art of Memory Routledge & Kegan Paul (also Pimlico 1992)

Zeki S 1974 ‘The mosaic organisation of the visual cortex in the monkey' in Bellairs R and Gray EG (eds) Essays on the Nervous System: a Festschrift for Professor JZ Young OUP, Oxford

Zeki S and Lamb N 1994 ‘The neurology of kinetic art' Brain (117) 3 pp607-636

Zucker SW and Davis S 1988 ‘Points and endpoints: A size/spacing constraint for dot grouping', Perception (17) 2 pp229-247

Internet Resources

Since it is an Internet-based standard, there are a number of resources in support of VRML to be found on the Internet itself. Indeed, due to its rapidly changing nature, this is usually the best place to find the current state of affairs. Much printed material is soon out of date.

- The VRML Repository (http://www.sdsc.edu/vrml), with a mirror site in Switzerland (http://www.vrml.ch), is a major site for information on the subject. It contains software, sample worlds, documentation and links to projects and worlds. Being semi-official, this site is likely to provide the most stable source of information on the subject and up-to-date links to other sites, and should be the first point of contact.

- For a more local flavour, there is a UK VR Special Interest Group - ukvrsig - (http://www.crg.cs.nott.ac.uk/ukvrsig/vr-sig.html) with local chapters in various parts of the country. In addition, there is the London Virtual Reality Group, LVRG (http://dougal.derby.ac.uk:80/lvrg/). These groups occasionally organise meetings where those new to VR and VRML can meet those with some experience. Both the UKVRSIG and LVRG also run mailing-lists where queries can be put, and are often answered.

- The magazine Wired' is supporting VRML in providing the VRML Forum' (http://vrml.wired.com/). There is also the vrml mailing list. This can be joined by sending mail to majordomo@wired.com (No subject field, message body to be: subscribe www-vrml your-email-address). This generates a lot of mail at present as the subject is expanding, so there is an alternative digest version. It can be joined by emailing majordomo@wired.com, with the message www-vrml-digest.

- A good library of pointers to VRML material is provided by Weblynx (http://weblynx.com.au/virtual.htm).

- Serch is a database of links to multimedia sites, especially VRML (Virtual Reality Modelling Language) URLs, from Aereal Inc. Serch can be accessed through a standard WWW browser, or a VRML browser in Virtual Reality. (http://www.virtpark.com/theme/cgi-bin/serch.html)

- There is an index at Mesh Mart (http://cedar.cic.net/~rtilmann/mm/vrml.htm). In particular they maintain the VRML Update, an online magazine' for vrml issues (http://cedar.cic.net/~rtilmann/mm/vrmlup.htm).

- There is an online tutorial for version 1 by Pete Clark at (http://www.mwc.edu/~pclark/vrmltut.html).