Virtual Reality for Animals:
linking concrete and abstract reasoning through action in virtual space

John A. Waterworth

Department of Informatics
Umeå University
S-901 87 UMEÅ
Sweden

jwworth@informatik.umu.se

Abstract

The key feature of VR (Virtual reality) is the way we can now concretise the abstract. The key question is how the abstract, the purely mental, relates to the physical, the concrete, the bodily. Can the experiential and cognitive aspects of mental life be catered for, indeed blended, through smooth human-computer interaction in a VR? In this speculative paper, I claim that they can be more than blended, that experiential interaction with concrete objects allows abstract, purely cognitive work to be bodily located within a virtually-physical landscape, and that this has to be designed in addition to the physical, bodily experiences which may (or may not) stand for abstract entities and relationships. So I suggest that there are three domains for the design of VR interaction: VR as pure bodily experience, VR as one or more embodiments of abstract knowledge, and VR as a link between the bodily and the abstract in the human participant. Including the third domain allows VRs to be designed with the express purpose of supporting cognition and creativity through virtual spatial action and experience. The approach draws loosely on the experiential theory of meaning of Lakoff and Johnson as well as Jung's views on the relationship between the conscious and the unconscious.

1 Introduction: VR for Animals?

Like all animals, people have bodies and a range of senses, perceptual and motor skills that enable them to function in the physical world. Increasingly, human-computer interaction (HCI) design aims to tap these sensori-motor skills so that users do not need to engage their cognitive capacities for abstract thought much in order to interact with the system (Waterworth, 1996a,b). Instead, they can function on the basis of the unconscious, automatic or automatised skills that underlie pattern recognition (Nygren et al., 1992) and smooth sensori-motor performance. Although the skills are unconscious, their results are not. An experiential account of meaning suggests that cognitive reasoning makes sense to us because we relate mental concepts to essentially bodily experiences (Lakoff and Johnson, 1980; Johnson, 1987; Lakoff, 1987).

In virtual reality (VR), HCI is everything, and in designing HCI we are often making solid the insubstantial, concretising the abstract. In VR, where everything is a thing, this is very clear, but VR is not fundamentally different in this respect from the general trend of developments in HCI. VR is thus the culmination of our recent, technologically-motivated, cultural progression towards the concrete and away from the abstract. Look at television advertising. Concepts which were once conveyed by a human expert (in a white coat) explaining the advanced science that had gone into the manufacture of a soap powder and the way in which the science works, are now communicated by apparently-3D animated characters representing both the wonderful qualities of the soap powder and its chemistry and the various different kinds of dirt and other stains that bring out the heroic qualities of the personified powder.

Now we create interactive virtual worlds where things which appear to be objects stand for abstract entities and processes. This means that we do not have to do mental work ourselves to make the abstract concrete and, hence, comprehensible. We are developing VR for animals (actually, of course, for people as animals), since animals share our sensori-motor skills but not our capacity for abstract thought. For the first time, we can be seen to be developing a style of information representation that could be applied to animals as readily as to people. A good test of a VR would thus seem to be whether it can convince and be used by a fairly smart mammal - say a dog, a cat or a pig.

People had techniques for using the concrete to render the abstract more manageable centuries before they had the technology to do so. Memorisation of large sets of information, for presentation in the oral tradition, was often achieved by deliberately and laboriously associating to-be-remembered items with locations in real, physical space - such as a well-known building or "memory theatre"(Yates, 1984) - the "method of loci". In some ways, we have moved full circle, and current anxieties about the "end of literacy" reflect this (Waterworth, 1992). The fear is that we are heading towards a new Dark Age of unreflective thought and immediate sensory gratification. This is reinforced by the fact that now the abstract to concrete mapping involves no conscious effort on our part. Might we become unaware that there is such a mapping?

Such concerns motivate our asking just how the concrete and the abstract can and should be integrated in virtual worlds. Another concern relates to the metaphoric nature of VR, given the way metaphors work through the tension between sameness and difference (see, for example, Snodgrass and Coyne, 1992). The object can stand for what it models, and aid comprehension, because of the similarity between the two, but also because of the difference between the two. This implies that any metaphorical representation is **partial** and, to an extent, **arbitrary**. What, then, is the status of understanding based on interaction with apparently-physical objects which stand for abstract ideas, but only partially and somewhat arbitrarily?

2 The Two Streams of Mental Life

As animals living in a physical world, but also capable of abstract thought, we function on a moment-by-moment basis by integrating two streams of information processing in consciousness: the sensori-motor and the conceptual (or "cognitive"). This section examines how we integrate these two worlds of experience - which are also the concrete and the abstract - as a step towards creating more robust models of HCI for VR.

Although sensori-motor functioning is largely automatic in its execution, its results are available to conscious inspection and it is under executive conscious control to a large extent. In contrast, conceptual processing is largely conscious in its execution but (counter-intuitively perhaps) is predominantly **not** under conscious control. For example, when we take our turn in a conversation, we generally do not know what we are going to say, and only find out when we hear ourselves speak. In fact, recent evidence suggests that conceptual processing is controlled through a degree of anchoring in planned sensori-motor processing (see, for example, Hendriks-Jansen, 1996).

Additional support for this view comes from studies of mental dysfunction (e.g. Sacks, 1985), which show that to lose the concrete is, psychologically speaking, far more damaging than to lose the abstract. A man who can only think abstractly cannot function in the world, whereas one who is locked in the concrete lives in the here and now, although his thoughts and corresponding reactions are simple (but not simpler than a dog or cat, and far more complex than a cockroach). In functional people, the concrete is primary and dominant, even in the most abstract of thinkers. Abstract thought "floats" on a bed of normal animal action and reaction. And so it should be with VR environments which seek to support, or at least not discourage, abstract thought ("reflective interaction", to use Norman's term; Norman, 1993). It is better to lose your mind (abstract reasoning) than your body (concrete reasoning).

What stands against the claim that human cognitive skills (as opposed to the perceptual, sensori-motor abilities we share with animals) are actually rooted in such animal qualities? In a word - and one that has had a profound influence on the development of HCI as a scientific discipline - "cognitivism". This is the idea that man functions, learns and understands through his ability to represent the world in a set of abstract symbols that can be consciously manipulated. In other words, man as a symbol processor, not an animal, based on what was the dominant metaphor of human nature throughout the 1970s and 1980s - mind as computer (see also Waterworth, 1996b).

But there are many problems with purely cognitivist accounts of man. As Schön (1987) has pointed out, to see man as purely a processor of information is to distort the learning problem by identifying knowledge with computational models. But models are a way of informally testing concepts; they are not themselves concepts. Having a model is not the same as understanding something. More fundamentally, cognitivism is rooted in the philosophy of objectivism (Johnson, 1987). This is the rather simple idea that the world is at it is, whatever any person believes or experiences it to be; human experience is seen as incidental to meaningful thought and reason.

Objectivism assumes that meaning is objective because it consists only in the relation between abstract symbols and things in the world. But, as Johnson points out, there are both logical and empirical problems with objectivism. Putnam (1981), for example, shows logically that the meaning of abstract symbols cannot reside purely in some direct correspondence between the symbols and the world. We may say that meaning only arises when a person experiences it. Meaning is thus prior to, and cannot only arise from, the manipulation of abstract symbols. Johnson's catalogue of empirical support for this view includes human categorisation behaviour (e.g. Rosch, 1975), concept framing in different cultures reflecting different experiences (e.g. Fillmore, 1975), metaphor, metonymy and polysemy (e.g. Ortony, 1979; Fauconnier, 1985), semantic differences over time, non-western conceptual systems, and the growth of knowledge over time (Kuhn, 1970).

The idea of two streams of mental activity is implicit in much existing work following the largely cognitivist HCI tradition. Generally, the assumption is that an interface should not be cognitively demanding to use, so that the maximum amount of cognition is available for other problem solving (Norman, 1986). This is the rationale for direct manipulation (Shneiderman, 1988) which, from the early 1980s, has been the dominant style of interaction (with graphical screen objects). However, this early adoption of the beneficial effects of tapping sensori-motor skills did not lead to an abandonment of cognitivist HCI design. On the contrary, the claim was that it freed mental capacity for more abstract problem solving at the interface.

But there is evidence that things don't always work this way. Golightly et al. (1996) cite studies showing that, sometimes, direct manipulation results in poorer problem solving, because the manipulation in some way inhibits cognitive work. Or to put it the other way, if one can manipulate things to produce a solution in a reasonable time period, one is less likely to put as much effort into solving the problem mentally (i.e. by abstract cognitive processing). The apparent paradox disappears when one considers the location of the problem to be solved. If the problem is embedded in the interface ("Puzzling Interfaces"; Golightly et al., 1996), then users will tend to solve it by manipulation rather than thought; and as a result, they tend to take more time to reach the solution. If the problem is independent of the interface, any mental effort spent on the interface would distract from problem solving, so a direct manipulation interface will reduce the time to reach a solution.

Where is the location of the problem to be solved in VR? If we take as an example an environment developed for students to learn Newtonian physics through direct interactions with highly manipulable objects and forces (Dede et al., 1996), is the problem embedded in the interface, or is the interface merely a distraction from the problem to be solved? In this case, and viewing things through cognitivist eyes, it seems the interface **is** the problem, and so we might expect a very concrete, manipulable interface to interfere with problem solving, by the argument given in the preceding paragraph. But the rationale of design in VR is different. It is not cognitivist, concerned with how limited symbol processing capacity is allocated. It is experientialist (Lund, 1996) and thus primarily concerned with producing experiences that will lead to desired kinds of perceptions in users. From this perspective the interface is not the problem, it is the solution. More time spent manipulating things at the interface gives more insights into the solution, not less time to spend thinking about the problem.

The work outlined above points to the inadequacies of purely cognitivist accounts of man as primarily a processor of symbolic information, and a resurgence of a more balanced view that acknowledges that we are first and, arguably, foremost animals. Much of our behaviour is automatic or automatised and it is rooted in bodily action and the way, as humans, we have to integrate physical existence with abstract thought.

Further evidence for the importance of a more "physical" conception of mind than is allowed by cognitivism comes from a few studies of people who earn their living dealing with abstract concepts, so-called "knowledge workers". It has been claimed that these paradigmatically cognitive workers often use physical space to create "holding patterns" for loosely-structured collections of ideas (e.g. Mander et al., 1992; Kidd, 1994; Marshall and Shipman, 1995; Waterworth, 1996a). This echoes the early mnemonic techniques of orators (Yates, 1984), but without the memorisation element. In this case, the holding pattern preserves a perceptual arrangement that has been observed but not memorised, perhaps because of the work involved or because this arrangement has not yet been understood sufficiently abstractly for it to be classified. There is even a suggestion that when it has been "understood" enough to be filed away, it is ready to be forgotten forever.

The "pattern holding" behaviour of knowledge workers shows us how the abstract may sometimes be anchored in the concrete in the service of improved cognition. Another way of looking at this is as the embedding of the cognitive in the sensori-motor, the bodily. So, not only is the body in the mind (Johnson, 1987), the mind is in the body is in the mind.

Several experiments have been designed and some are currently in progress, to test these ideas. At the time of writing, only preliminary results are available, but these are highly suggestive and support the basic idea that cognitive work can be supported by anchoring it in concrete reasoning skills (which is not a new finding). Further, those concrete reasoning skills can be enhanced by the use of three-dimensional spatial VR (which is new).

3 Routes and Places

3.1 Space as Procedure versus Declaration

People act naturally in spatial worlds, finding their way around by using sophisticated sensori-motor skills. Spatial skills are deeply encoded in the human mind and are predominantly perceptual rather than cognitive in nature. These skills are amenable to automatisation, which means that they come to require little or no conscious control. Once automatised, these types of skills can be performed with very low error levels. The necessary knowledge is said to be stored in motor memory, and is accessible for (largely unconscious) use, but not for (conscious) recall. But space also has meaning for people, and they find spatial locations intrinsically memorable (recallable).

This dichotomy can be explained by the difference between two fundamentally different ways in which people relate to space: as marking aspects of a procedure for travel, and as a declaration of the differences between things. In the first case, generally known as following a route, the knowledge is largely unconscious. In the second, more commonly referred to as putting things in their place, the knowledge is conscious.

Mnemotechnic methods link route following with location (and the items "stored" there) retrieval. By mentally following a route, to-be-remembered items are found and recalled. The important thing for these methods to work is that the space is familiar to the person using the method. This is personal to the user and suggests the idea of personal spaces (see 3.2, below), in contrast to the public places we happen across when following a sequential path.

It seems plausible that virtually-spatial interfaces to information systems can themselves serve as artificial memory for their users. Because such systems are artificial but not merely imaginary, they can capitalise on both imagery and spatial memory in the recollection of the location of items of information. We might thus expect that once users have explored the 3D objects and seen what items are associated with which location in the structure, they would experience little difficulty in recalling where to find a particular item the next time it were needed. This has interesting and significant implications for the design of HCI in general, and especially for the design of VR environments that support conceptual work through interaction with virtually-concrete objects and forces.

3.2 The Mnemonics of Space

This section describes work planned and in progress to investigate the role of spatial action in mediating cognitive reasoning. This is a research agenda, with predicted outcomes. My purpose in presenting this before the full results are available is to give a feel for the kinds of phenomena entailed by exploring, with virtual spatial action, the links between physical and mental reasoning.

3.2.1 Spatial versus Non-spatial Bookmarking

One environment where spatial interaction has received considerable attention is that of the World Wide Web. Given the incomprehensible structure of the Web, there is a need for users to maintain a personal collection of Web locations, usually catered for by some kind of bookmarking functionality. However, collections of bookmarks bring their own problems, and Web browsers typically offer some way of organising a user's collection of bookmarks to help prevent users getting lost within their own bookmark lists. Methods for achieving this include sets of folders and other hierarchical filing methods. Another approach is to allow users to spatialise the collection, creating what is referred to as a "Personal Information Space" or just "Personal Space"(Waterworth, 1996a).

A survey of several regular users of the Web is currently in progress, focusing on how they structure bookmarks, and on what mechanisms they use for such structuring. An experiment was initiated to compare the use of spatial representations of bookmarks with non-spatial. An important aspect of such a study is that the data should be captured over a realistically long period of use (of the order of several months) and should include as many representative users as possible.

Although the final results are not available at the time of writing, our prediction is that spatial bookmarking will result in better recall of earlier interactive behaviour. What this means is that spatial bookmarks result in fewer problems remembering whether a location has already been bookmarked (and thus less duplication of bookmarks) and generally better recall of earlier interactive sessions and the work that was completed during those sessions - the task context.

The important point here is that the spatial structure serves as a framework for remembering the abstract, reflective processing carried out by the user during its construction and later use, not merely as a filing system for interesting Web pages.

3.2.2 Using Locations as Memory Cues

Another study was designed to examine the effects of a fixed versus a changing spatial landscape on plan memorisation and recall. The other main variable to be manipulated is whether the information is stored at locations in the landscape or not. So there are four main conditions, as shown in Figure 1, below.

	Same spatial arrangement	Different spatial arrangement
Information in location	A	B
Information not in location	C	D

Figure 1 - Experimental conditions of information retrieval, space and locations

We would expect that when information is stored at specific locations, recall is disturbed when we rearrange the spatial landscape. In other words, we would expect better plan retrieval in condition A then condition B, and this seems to be the case.

But what of conditions C and D? The surprising finding is that plan retrieval may also be better for condition C than condition D (though not as good as in condition A). In other words, changing the spatial landscape could affect the retrieval of plans even if those plans are not associated with locations in the landscape[1]. How can this be?

[1]An example might make this clearer. Consider planning a journey where different activities are to be carried out at locations along the route. In one case, the activities depend on the location - swimming in the sea, visiting a shrine, exploring a city, say. In the other, they do not - writing the first chapter of a book at location 1, then driving on to location 2 and writing the next chapter, and so on,.

The suggestion is that humans are able to anchor information mentally in locations, even if those locations are not actually used for information storage. This is because we must navigate whatever landscape we are presented with, and we are built to do this efficiently and largely unconsciously. In the process, we are also always engaged in cognitive, abstract thought which may or may not be physically (or virtually-physically) related to the landscape in which we find ourselves. In either case, we tend to anchor our cognitive stream of processing in our physical route-following. We can be said to control our conscious but not easily controllable abstractions with our largely unconscious, but controllable, physical (or virtually-physical) actions.

The next study was designed to emphasise this point and to assess the importance of real (or realistic) action within physical space to maximise the effect.

3.3 Cognitive Plans and Spatial Navigation

If the abstract is somehow anchored in the concrete, how important is realism and three-dimensionality to this effect? To address this question, three different versions of the same landscape were created. In one, real physical space (a set of interconnecting rooms) is used. In the second, a reduced scale model of the same structure was constructed out of cardboard. In the third, a two-dimensional map of the same layout is provided.

Again, the task to be performed by our volunteers is that of plan memorisation and later recall. Performance is tested by scoring how accurately plans can be reconstructed, with the landscape representation as cue. We predicted that the best performance would be produced by the real, linked rooms condition, but that there would be little difference between the cardboard model condition and that of the sketch. Note that this is the opposite of what one would expect from studies of navigation in hypermedia, for example. There, high level overviews generally improve way-finding as compared to simply following links (Wright, 1991). The difference lies in the fact that we are not assessing how well people can find their way - the way is supplied - but rather how well they remember "the thoughts they have along the way", as it were. Preliminary results are compatible with these predictions. Figure 2 shows a high-level overview and a street-level view of one landscape.

Since we have confounded degree of physicality with realism in this study we are also trying to separate them out by reconstructing similar conditions on screen, varying only degree of realism. The surprising finding this time would be that the degree of realism has little or no effect. This is compatible with the idea that it was, rather, the closeness to navigation in real physical space that was important; and this is where VR comes into its own.

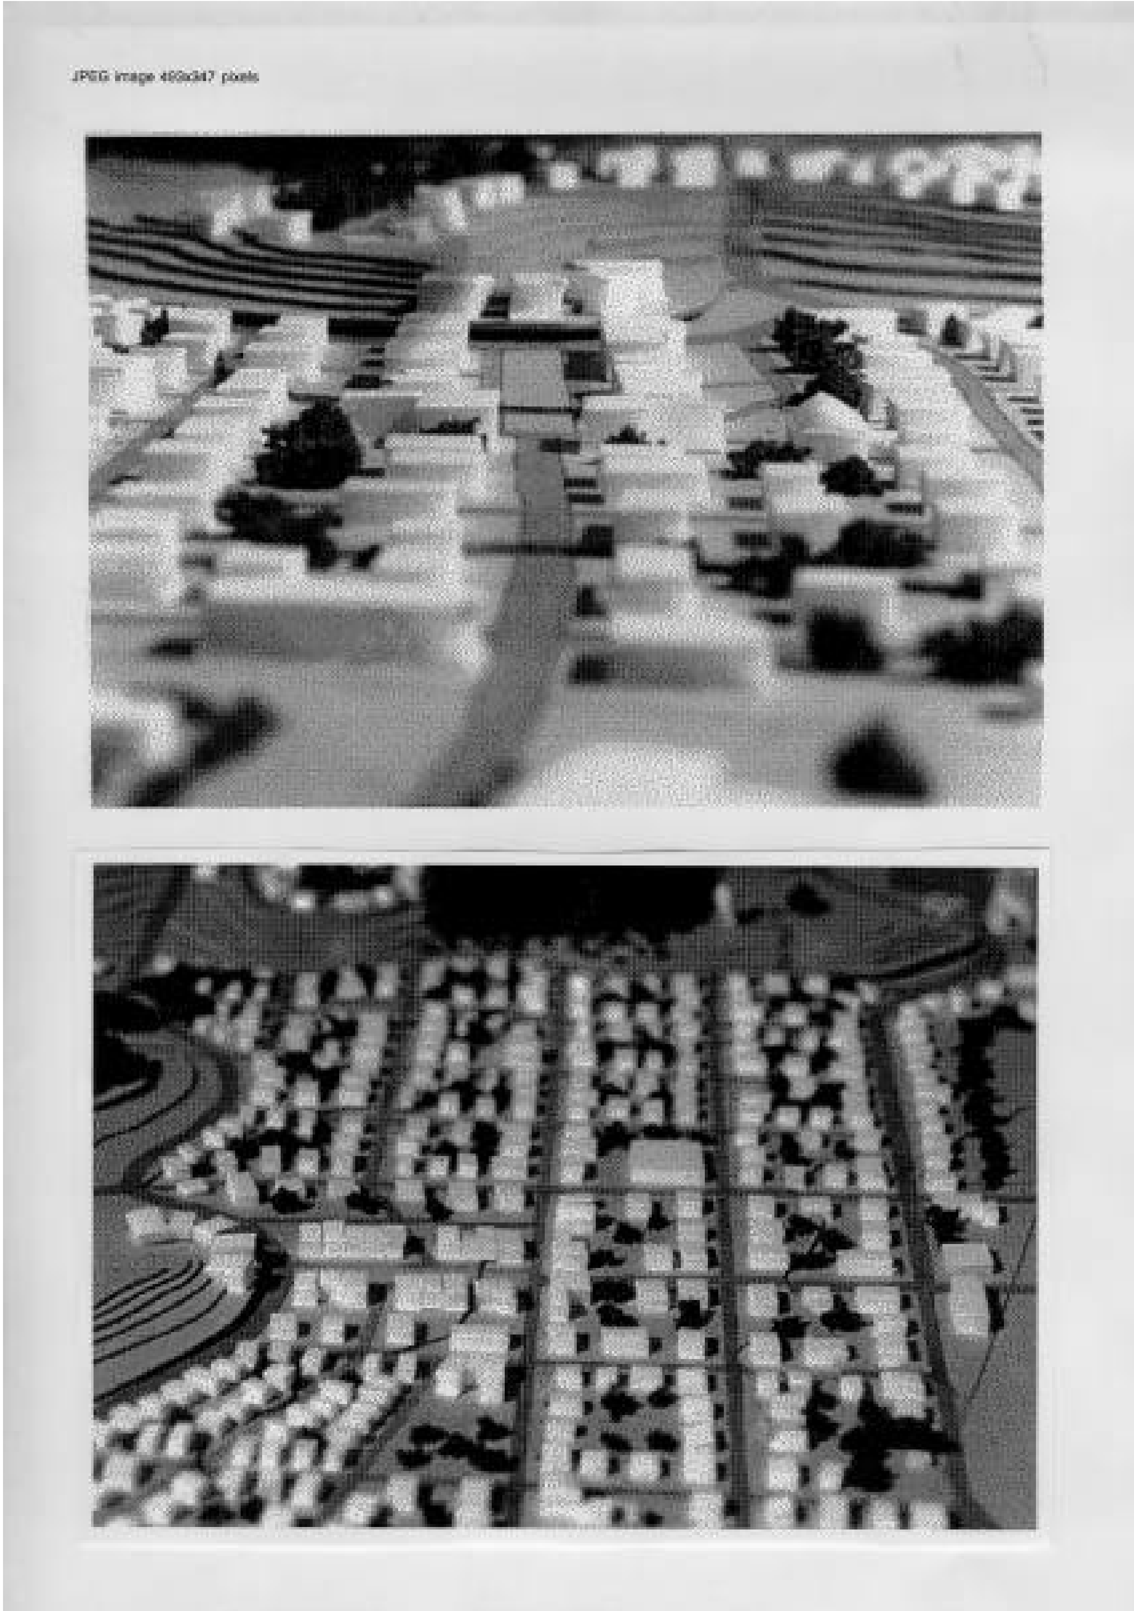

Figure 2 - A High-Level Overview and Street-Level View of the Same Landscape

4 Linking the Streams

> The moment of truth, the sudden emergence of a new insight, is an act of intuition. Such intuitions give the appearance of miraculous flashes, or short-circuits of reasoning. In fact they may be likened to an immersed chain, or which only the beginning and end are visible above the surface of consciousness. The diver vanishes at one end of the chain and comes up at the other end, guided by invisible links.
>
> A Koestler, 'The Act of Creation'

4.1 The Three Domains of VR Design

We can see now that there are actually three domains of design for the HCI of VR. Firstly, VR as pure bodily experience; secondly, VR as one or more embodiments of abstract knowledge; and thirdly, VR as a link between the bodily and the abstract. Work in VR design has focused particularly on the first of these; this is the standard model of VR for most existing applications, such as games, medical visualisation/modelling and architectural simulations for virtual 'walk-throughs'.

The second category of design, where virtual objects and events stand for more abstract concepts with which it is hoped a learner will become familiar through the process of interaction, is less widespread but is typical of pedagogical applications of VR such as Newton's World (Dede et al., 1996). In this paradigm, VR is seen as an extension of multimedia instructional techniques. The learner experiences the interface, and in so doing is expected to reflect on the concepts for which the objects and events stand. Experiential, bodily interaction is intended to support cognitive, reflective thought.

Doubts concerning this approach revolve around the issue of whether interaction with a partial, arbitrary way of representing the concepts to be grasped results in an abstract and generalisable understanding of those concepts. The argument in favour of such an approach is generally that any understanding is better than none, and/or that a partial understanding can develop into a more accurate one through a process of concept displacement (Schön, 1963).

The focus of this paper is the third realm of design; where a virtual reality is designed to serve as the link between the bodily, the experiential and abstract, reflective cognition. In this approach, a VR is designed to integrate the two streams of human thought, in a similar way to that in which the mnemonic techniques of the orators supported extended performances in the oral tradition. The important difference, of course, is that the VR does the work of association for us.

4.2 Opening the Door to Creativity

In a recent paper (Waterworth, 1996b) it was claimed that one way in which information technology can support human creativity is through sensory enhancement, by the development of "synaesthetic media". Such systems allow the human participant to choose the modality (or modalities) through which information is experienced. So, for example, data on stock market trends could be presented as sound - a financial symphony, as it were - rather than as the more usual tables and graphical charts. In this way, the naturally-occurring but abnormal phenomena of synaesthesia, such as "hearing" normally-visual sights or "touching" normally-heard sounds can be duplicated artificially.

Many natural synaesthetes are creative individuals, most commonly in the artistic sphere. The further claim was that such synaesthetic media can improve human creativity by allowing the same information, the same aspect of reality, to be experienced in a multiplicity of ways. Since we don't know what reality is, the more ways we experience it, the more likely we are to have novel insights. An assumption of this approach is that we normally filter out such rich experience to increase our immediate chances of survival. Creativity by this view rests on sensuality, and is a luxury that we can only afford when day-to-day survival is reasonably well assured.

However, there is more to creativity than sheer sensuality. Creative insights have often been associated with the activities of the unconscious mind. Perhaps the most familiar example is that of breakthroughs following the experience of a particularly striking dream. If we take a Jungian perspective, we can account for this dimension of creativity in terms of the relationship between the conscious ego, the individual unconscious, and the collective unconscious (Jung, 1953-1979a). More prosaically, emotional reactions betray unconscious processes:

"In 3D, the force of what you're looking at is much stronger. It's almost an emotional impact. Even though you know perfectly well intellectually that the image is sitting there on the screen, if you swing around, people back up so as not to get hit by it. That suggests that the information is being interpreted in a much more powerful way. You are engaged with it, not just interpreting it." - James Bower, neuro-biologist (quoted in New Scientist, 1994).

Linking the cognitive, abstract stream of thought with the physical, bodily stream amounts to linking the conscious with the unconscious, since most bodily processing is unconscious.

The ability to create such links explicitly (in the third realm of design) also suggests that we might take the idea of the collective unconscious seriously, and accept some kind of relationship between it and the physical world. From the Jungian viewpoint, such a relationship may potentially become conscious, and to support this, we should provide for casual, not causal relations between the two. Jung's term for this kind of acausal connection was "synchronicity" (Jung, 1953-1979b). Such synchronicitous events often form the basis for literature, drama, cinema and, arguably, the visual arts and religion. If human meaning is not a function of causal connections but of bodily experiences (Johnson, 1987), we can see that these may be real, imagined or dreamed, arising from the human sensori-motor system (the unconscious) and the world of collective myths and spirituality (the collective unconscious).

To open the door to creativity, then, we need to design for the human participant to be conscious of synchronicitous links between the mental and the physical. We also need to provide as many alternative concrete representations of abstract ideas as possible, since each is partial and to an extent arbitrary. And we may wish to experiment with populating a VR with mythic objects and landscapes that are likely to encourage conscious access to the world of the collective unconscious.

4.3 The Location of Consciousness

One way of characterising the dichotomy between the two streams of mental life is as a question of where consciousness is located. Although physical processing is largely unconscious it is not, of course, entirely so. We are aware of what we are feeling and doing. And while cognitive processing is largely conscious it is not entirely so - we are often unaware of where our thoughts come form or are going. The problem of consciousness is that of limited capacity. We have very little attention at our disposal and we must share it between sampling from the physical environment (and controlling physical actions) and carrying out conscious mental work (or cognition). In this way we switch attention between the mental and the physical.

When our conscious processing load is heavy, our experience of duration is short - "time passes quickly" (Waterworth, 1983). We pay little attention to our bodies or the world around us. And when our conscious processing load is light, duration seems long - "time passes slowly". Figure 3 illustrates this graphically.

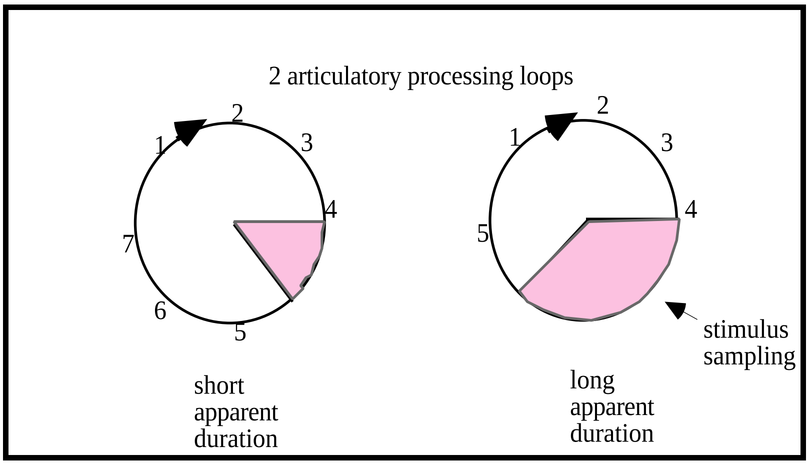

Figure 3 - Two Articulatory Processing Loops Varying In Cognitive Load

We can see the non-processing stimulus sampling zones of the loops as windows on reality. I would also argue that they provide the opportunity for creative insights to enter consciousness. In VR design, we can manipulate both experienced space and experienced time, and so encourage creativity (if that is our design aim).

Anchoring the mental in the physical reduces the demands on our limited attention span, our consciousness. Loss of the mental does not (usually) result in our losing track of planned physical actions. But if we lose track of our mental stream, re-enactment of the most recent physical sequence of events enables us to recover the mental sequence[2]. The physical landscape contains and supports cognition.

4.4 The Future: Landscapes for the Creative Mind

The role of the physical (or virtually-physical) landscape in supporting cognition points to the importance of the actual characteristics of particular landscapes in VR design. Figures 4-6 suggest a range of landscapes, each of which might be expected to have effects beyond the immediate conscious impression.

An example of the application of this kind of approach to the design of interactive environments is that of Aguilera (1997). But here, the possible alternative paths the participant may follow are determined by the designer, who takes on the role of Creator.

An agenda for future research includes the idea that the participants should themselves have access to the creative process, and hence the need not only to design mechanisms for linking the abstract and the concrete but also tools for participants to experiment with ways of linking.

---

[2]A familiar example of this is when a good idea occurs to one while, say, making a cup of coffee in the kitchen. Wandering into the living room to drink the beverage, one finds the idea has been lost, seemingly irretrievably. However, if the physical movements carried out in the kitchen are re-enacted, the idea will often be recalled.

Figure 4

Figure 5

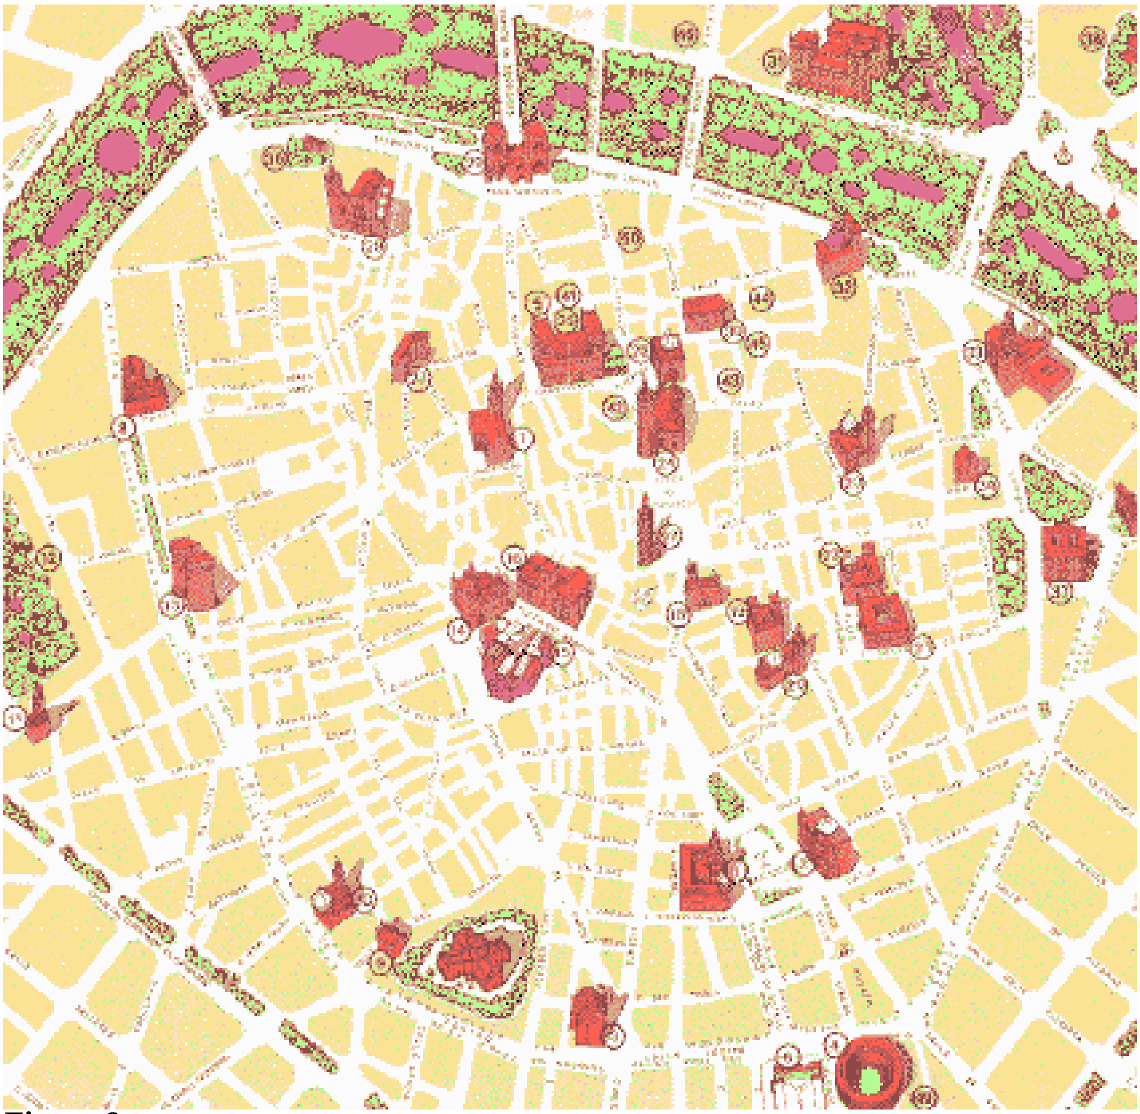

Figure 6

Figures 4-6 Alternative Virtual Landscapes

Imagine a VR where the user is free to explore a set of apparently-concrete objects, and to deconstruct and reconstruct landscapes of their own from such objects. They are also able to assign links between "environment-objects" and "information-objects" - to put information in places according to their own preferences, in the way we put food-objects in kitchen cupboard-objects at home. This would be a basic environment for using concrete reasoning skills to support abstract reasoning.

To go further and support the creative process more directly, we can go at least two stages further. Firstly, users would be allowed to chose the modalities through which objects are experienced. Secondly, we might immerse users in symbology, identified from mythology, and allow objects to provide bodily access to ideas and forces from the individual and collective unconscious.

5 Conclusions

Your mind is in your body, and your body is in your mind.

Your mind is in your body, because you anchor cognition in bodily interaction with the physical world about you. Your body is in your mind, because to make sense, to have meaning, each rational concept ultimately derives from bodily experiences - actions and feelings - via images of such experiences (Johnson, 1987).

The idea of VR for animals is facetious, but is intended to draw attention to our cultural progression away from the abstract and towards the concrete, a progression exemplified by developments in virtual reality. If we were to limit our VR design efforts, as we currently do, to the construction of purely bodily experiences, we can legitimately see this effort as that of designing for people as animals (or even for animals themselves). In some spheres, such as medical and other applications calling for a high degree of dexterity, this is appropriate (but might limit applicability to apes with opposing thumbs).

Although we certainly have cognitive capacities that distinguish us from animals, we are still the same genetic organisms that used memory theatres. Our cognitive skills have not improved. Our needs, in terms of purely "cognitive artifacts" (Norman, 1993) are very modest. Such artifacts are designed to support "reflective interaction" (i.e. cognitive) as something apart from and antagonistic to experiential interaction. Contrary to that approach, however, we are more likely to improve our reasoning skills, and become more creative (see also Waterworth 1996b), if our senses and bodily experiences are addressed in the process of designing concrete representations of abstract entities and of explicit links between the two. Explicit abstract-to-physical and physical-to-abstract linking (not merely matching), and tools for manipulating objects and linkages can provide a VR environment that serves not only either cerebral humans or physical animals, but the human animals we are.

VR can support not only concrete but also abstract reasoning skills. It can be used to improve thinking, and not only of concepts that can be readily embodied as physical entities. But, it can only do this insofar as VR designers provide environments in which abstract reasoning can be embedded (but not embodied) in concrete reasoning skills whose execution is highly automated in the sensori-motor systems of people. This is the normal way in which people anchor and control abstract reasoning by attachment to sensori-motor plan execution. People can then use a suitably-designed VR environment to enhance these skills, because the environment itself can do the anchoring work.

Designing such VR environments remains the major challenge, and this paper has not shown how to go about it, only suggested that it is possible. The partial and arbitrary nature of any VR representation is no more of a stumbling block than any designer faces; any artifact can be designed in a multiplicity of ways and none of them will be satisfactory in all ways. But in VR at least we are free to provide as many views of reality as we choose. Each view addresses our animal nature, but multiple views, in multiple modalities, provide food for human creativity.

Virtual reality may serve as a means of reconciling our animal and rational natures. When we show these multiple aspects of a reality that cannot be experienced in its totality by animals, shifting senses and views in multiple ways, what will happen? We don't know. But then we don't know what reality is. The overall effect of VR is the revelation of more and more reality. This is not so surprising: even animals dream.

5 Acknowledgements

Thanks are due to Kristo Ivanov, Torsten Nilsson and Thomas Ahlmark for very valuable suggestions for improvements to the paper (not all of which I have adopted in this version), and to Paul Roberts and Bessie the Cat for encouragement and inspiration.

6 References

Aguilera, J (1997) - 'Visual Poem/Virtual Poem'. This volume.

Dede, C, Salzman, M and Loftin, R B (1996) - 'The Development of a Virtual World for Learning Newtonian Mechanics'. In Brusilovsky, P, Kommers, P and Streitz, N (eds) *'Multimedia, Hypermedia and Virtual Reality'*. Springer Verlag Lecture Notes in Computer Science, pages 87-106.

Fauconnier, G (1985) - *'Mental Spaces'*. Cambridge, USA: MIT Press.

Hendriks-Jansen, H (1996) - *'Catching Ourselves in the Act'*. Cambridge, USA: MIT Press.

Golightly, D, Gilmore, D and Churchill, E (1996) - 'Puzzling Interfaces: the relationship between manipulation and problem solving'. *Adjunct Proceedings, HCI'96 Conference of the BCS HCI Group,* Imperial College London, August 1996.

Johnson, M (1987) - *'The Body in the Mind'*. Chicago: University of Chicago Press.

Jung, C G (1953-1979a) - *'Relations between the Ego and the Unconscious'*. From Two Essays on Analytical Psychology, Collected Works Volume 7, Second Essay, Paragraphs 202-295.

Jung C G (1953-1979b) - '*On Synchronicity'*. From The Structure and Dynamics of the Psyche, Collected Works Volume 8, Paragraphs 969-997.

Kidd, A L (1994) - 'The Marks are on the Knowledge Worker'. In *Proceedings of CHI'94 Conference on Human Factors in Information Systems* (Boston, April 1994), 186-191. New York: ACM.

Kuhn, T (1970) - *'The Structure of Scientific Revolutions,* 2nd ed.'. Chicago: University of Chicago Press.

Lakoff, G (1987) - *'Women, Fire and Dangerous Things'*. Chicago: University of Chicago Press.

Lakoff, G and Johnson, M (1980) - *'Metaphors We Live By'*. Chicago: University of Chicago Press.

Lund, A (1996) - *'Embodied Interfaces: Towards an experientialist approach to user interface design'*. Working paper: Department of Informatics, Umeå University, Sweden.

Mander, R, Salomon, G and Wong, Y (1992) - 'A Pile Metaphor for Supporting Casual Organisation of Information. In *Proceedings of CHI'92 Conference on Human Factors in Information Systems* (Monterey, April 1992), 627-634. New York: ACM.

Marshall, C C and Shipman, F M (1995) - 'Spatial Hypertext: Designing for Change'. *Communications of the ACM, 38* (8). 186-191.

New Scientist (1994) - 'Fantastic Voyage into the Virtual', written by Bob Holmes. *New Scientist,* 2 July 1994, 26-29.

Norman, D A (1993) - *'Things that Make Us Smart'.* Reading, Mass: Addison-Wesley.

Nygren, E, Lind, M, Johnson, M and Sandblad, B (1992) - ‘The Art of the Obvious’. *Proceedings of ACM CHI’92 Conference on Human Factors in Computing Systems* (Monterey, May 1992). New York: ACM Press.

Ortony, A (ed.)(1979) - *'Metaphor and Thought'*. Cambridge, UK: Cambridge University Press.

Putnam, H (1981) - *'Reason, Truth and History'*. Cambridge, UK: Cambridge University Press.

Rosch, E (1975) - *'Cognitive Reference Points'.* Cognitive Psychology, 7, 532-547.

Sacks, O (1985) - *'The Man Who Mistook His Wife for a Hat'.* London: Duckworth.

Schön, D A (1963) - *'Displacement of Concepts'.* New York: Humanities Press.

Schön, D A (1987) - '*Educating the Reflective Practitioner'.* San Francisco: Jossey-Bass.

Snodgrass, A and Coyne, R (1992) - 'Models, Metaphors and the Hermeneutics of Designing'. *Design Issues, 9,* 56-74.

Waterworth, J A (1983) - *'The Influence of Variations in Cognitive Processing on the Perception of Time'*. PhD thesis, University of Hertfordshire, UK.

Waterworth, J A (1985) - 'Memory Mechanisms and the Psychophysical Scaling of Duration', *Perception, 14,* 81-92.

Waterworth, J A (1992) - *'Multimedia Interaction: human factors aspects'*. Chichester, UK: Ellis Horwood (Simon and Schuster International).

Waterworth, J A (1996a) - 'Personal Spaces: 3D Spatial Worlds for Information Exploration, Organisation and Communication'. *Proceedings of BCS Conference on 3D and Multimedia on the Internet, WWW and Networks,* Bradford, UK, 17-18 April 1996.

Waterworth, J A (1996b) - 'Creativity and Sensation: The Case for Synaesthetic Media'. *Proceedings of "Creativity and Cognition 96"*, LUTCHI Research Centre, Loughborough University, 29-30th April, 1996.

Wright, P (1991) - 'Cognitive Overheads and Prostheses: Some Issues in Evaluating Hypertexts'. In *Proceedings of Hypertext '91*, 1-12. New York: ACM.

Yates, F. A. (1984) - *'The Art of Memory'.* London: Routledge and Kegan Paul. First published in 1966.