EDUCATION AS EXPLORATION: BEING, GOING AND CHANGING THE WORLD

John A. Waterworth and Eva L. Waterworth, jwwaterworth@informatik.umu.se, Dept of Informatics, Umeå University, SE-901 87 Umeå, Sweden

Abstract

The rendering of abstract information as objects and events in apparently-concrete three-dimensional virtual worlds has profound implications for the way people relate to information. The importance of these developments for education are obvious, but the benefits are not so straightforward. The hallmark of virtual reality is the sense of presence, the feeling that one is actually there, within the virtual world. We discuss aspects of feeling present, including time perception and bodily responses, and their relation to conceptual versus perceptual information processing. In conceptual learning we typically develop an internal conscious model of relevant aspects of the external world. In perceptual learning we interact with the external world via increasingly unconscious skills.

Several authors have pointed to the importance of navigational features in virtual worlds, to allow people to use their real world navigational skills when dealing with virtual worlds of information. We present results from comparative studies of students learning to find their way around virtual worlds of varying closeness to the physical world. One finding was that students preferred the more object-based and less textual worlds, and felt a greater sense of presence there, but learning was no better. In another study, hypertext proved much more effective than an equivalent virtual world, though the sense of presence in hypertext was significantly less.

What is important in learning depends on the kind of knowledge to be acquired and what the student intends to do with it. Often, our evaluations of new modes of information presentation presuppose explicit, conscious, reportable knowledge as the goal. This may bias our interpretation of results in favour of entrenched technologies. On the other hand, exploration is not necessarily the best way to learn about unknown territory; reading a good guidebook may well make any excursions more worthwhile. In educational settings, we need to deliberately control the sense of presence to facilitate both vivid experiences and generalisable conceptual learning.

1 Introduction

One way of characterising education is as change in the balance between (and content of) conscious and unconscious knowledge. In any learning situation, students move through cycles of conscious access to information followed by unconscious utilisation of that information - often in the service of the next phase of conscious information access. Recently developed information technology, such as interactive multimedia and virtual reality, have a profound impact on this balance between the conscious and the unconscious. Often, they engender a strong feeling of presence, of being an active part of the reality that is portrayed. It is this feeling of presence that differentiates the new learning technologies from the old; and it is this same feeling of presence which is both their greatest strength and their most severe weakness.

In Section 2, we discuss the relationship between the sense of presence (of "conscious being") typical of interacting with new technologies and the learning process. After that, Section 3 expands on the idea of information access in learning situations as exploration of an information landscape ("going"). Section 4 emphasises the importance of students being active in achieving tasks in the virtual realities produced by new technology ("doing"). This includes not only exploring the terrain or searching for specific features, but changing the landscape and experimenting with the worlds in which they are virtually present. In Section 5 we revisit learning and education and also make a brief comparison of presence and absence before concluding, unsurprisingly, that education involves more than exploration. In the remainder of this introduction, we present some theoretical background to our views on learning and education.

1.1 Knowledge Representation and Organisation

One way of defining knowledge is as storing and organisation of information in human memory. Knowledge is organised information, which is part of a system, network or structured information.

A major debate in cognitive science has been about how knowledge is organised and used in memory. A characteristic that all approaches have in common is that knowledge is highly organised. One theory builds on the assumption that knowledge consists of numerous schemata, which facilitate our understanding of everyday events, and are based on previous experience. People develop schemata by repeatedly doing the same set of actions in a given setting. Another type of theory, which has been very influential approaches, asserts that knowledge is stored as a form of network with nodes and links. Information that has something in common is linked in some way. This way of viewing information is similar to the way a computer's memory is organised. Two of the theories common in the field of HCI are ACT* (Anderson 1983) and SOAR (Laird, Newell, Rosenbloom 1987), which both purport to describe how knowledge is stored in memory. Act* and Soar have some similarities and some differences. One characteristic that both theories share is the division of memory and production memory, while under SOAR there is only one long-term memory, a production memory, which is used for both declarative and procedural knowledge. SOAR could be said to

take an object-oriented view of the stored information, because both declarative and procedural knowledge are merged into the same chunks (Newell, Rosenbloom, Laird 1989). Network theories assumes that information exists in memory as independent units. These units are connected through links in an hierarchical network. An experiment by Tulving and Pearlstone (1966) demonstrates that information is not lost from long term memory, rather that it needs to have the right context in order to be retrieved.

There are generally considered to be three main types of knowledge stored in memory as analogical, propositional and distributed representations (Figure 1). Analogical representations are picture-like images, whereas propositional representations are abstract language-like statements that make presuppositions; as, for example, the car is in the garage. Distributed representations are a network where the knowledge is in the connections between the nodes. Analogical representations and propositional representations are regarded as symbolic representations while distributed representations are considered to be sub-symbolic representations.



Figure 1: Different types of knowledge representation in memory

1.2 Acquiring Skills: Novices and Experts

Allwood (1989) differentiates between two types of knowledge, skill knowledge and conceptual knowledge. Skill knowledge has also been called "knowing how" or "procedural knowledge", whereas conceptual knowledge is sometimes known as "knowing that" or "declarative knowledge". The development of skill knowledge can be described in accordance with Dreyfus and Dreyfus's (1986) five stages of skill acquisition (Figure 2).



Figure 2: Five stages of skill acquisition (Dreyfus & Dreyfus 1986)

In the first stage of acquisition novices try to learn simple objective facts and simple rules which are context-free. Experts on the other hand have knowledge of more facts and rules and also more sophisticated facts and rules. In addition to that experts' knowledge is more sensitive to context. Beginning learners have to concentrate when using their knowledge, while experts use this kind of knowledge automatically. Allwood (1989) claims that the knowledge a novice uses (when interacting with a computer for the first time) is of a descriptive, declarative character. Rasmussen (1986) describes a model of three levels of control of human actions (Figure 3) where he claims that the expert has skill-based behaviour and that this behaviour takes place at a more or less unconscious level. Learners, on the other hand, use knowledge-based behaviour, which is at a conscious level. Complete novices and experts can be said to be in a state of conscious being; they are present with the phenomena with which they are dealing. Active learners, on the other hand, are in a state of conscious doing; they are absent-mindedly unaware of major aspects of the phenomena they are dealing, because they are consciously dealing with mental representations of those phenomena (see later sections for more on presence and the conscious being-doing distinction).



Figure 3: Three levels of control of human actions (Rasmussen 1986 page 101)

Norman (1993) differentiates between two kinds of cognition, experiential and reflective. Experiential learning occurs when we perceive and react to events efficiently and effortlessly, while reflective mode is when we have to reflect, think and make decisions, which leads to new ideas and novel responses. But as Norman points out these two modes are not completely independent nor do they capture all of our thought processes. Reflection is more difficult than the experiential mode because it requires conscious work based on some mental structure and organisation. Experiential functioning on the other hand can be practised simply by experiencing it, and it is often enjoyable. It is characteristic of the way experts work, but also - somewhat paradoxically - of complete novices. This is because neither novices nor real experts they have internalised this structure so well that they can use it "automatically". In the case of novices, they have not yet even identified the necessary structure. In both cases, however, a strong sense of presence is likely to reflect the lack of conscious attention to the internal organising structure of knowledge. It is a characteristic of the active learner - neither a complete novice nor an expert - that they are not really present with the phenomena with which they are dealing.

The learning situation can be roughly separated in two extremes, traditional learning and learning by doing (Figure 4). This should only be seen as two extremes on a spectrum of learning situation. Traditional learning is when the learner is fed with selected knowledge and could be viewed as a passive process of the learner. At this extreme the view is that knowledge is possible to formalise. An example of this side of the spectrum, but not necessary at the extreme, is most of today's higher education. Which still emphasises theoretical learning and abstract thinking. At the other extreme, learning by doing, could be exemplified by learning handicraft, for example how a journeyman tailor during the 19:th century learned to become a tailor.

Traditional learning	Learning by doing
passive learning	active learning
a teacher lectures	learning by practice or from colleges
concentrates on important issues	treats all issues alike and it could hard to notice the important factors
ignores contexts	context is important
emphasise theoretical issues	emphasise practically issues
do not consider tacit knowledge or context	consider tacit knowledge and context
formalised knowledge	both formalised and tacit knowledge

Figure 4: A spectrum of learning situations

2 Being Present

The notion of presence refers to the extent to which consciousness is engaged in the currently perceivable situation - the concrete here and now - rather than in dealing with abstract conceptual representations of what is physically absent - the there and then. Conceptual learning is all about the generalised (or abstract) there and then. Perceptual learning is all about the generalised (or abstract) there and then. Perceptual learning is all about the here and now. Typically, both types of learning are initially conscious. And typically, perceptual learning is more rapidly, and less problematically, automatised into the unconscious. We are quite comfortable with the idea that perceptual skills - such as responding correctly to the colour of a traffic light - should become automatic and not require conscious consideration. And in many fields, such as sports, musical performance, fighter piloting, and car racing, only unconscious responses to the present situation are fast and accurate enough for optimal performance.

New technologies intrinsically emphasise perceptual presence - they tend to focus attention on the situation presented (present) rather than on concepts that might explain generalities or which linguistically describe situations which are not presented perceptually (i.e. are absent). Culturally, we are being driven by the technology we have developed towards the concrete (or rather, virtually concrete) and away from the abstract. This can be seen as a shift from mental problem solving - the conscious doing that underlies mental reflection - and towards virtually-physical manipulation - in the state of conscious being that underlies a sense of presence.

Below we roughly sketch an historical view of the effect of technological development on the way in which we think about consciousness, and the contents of consciousness. This will set the discussion of presence in a context that explains current concerns about the impact of new technologies on education.

Figure 5 presents a very crude time-line of the balance between consciously being and conscious doing (numerals in the following refer to the four major divisions shown in Figure 5).



Figure 5: A Time Line of Conscious Being (concrete reasoning) and Conscious Doing (abstract reasoning)

1. Pre-literate man - man before he really was man - was much like any other mammal. He responded to the immediate situation according to his perceptions of the current situation. His response tended to be automatic - instinctive - and unreflective. In other words, his cognition was unconscious. He did a lot of work with his legs and hands and almost none with his consciousness. But still, he was conscious. He did experience pleasure and pain, comfort and discomfort, hunger and fullness. And at some point, he started to do a bit of work, consciously. He started to plan his day to maximise his success in hunting and minimise his tiredness. He began to work out how to avoid dangers before he encountered them.

2. The development of language co-occurred with, and enabled, increased cooperation between early people in activities such as food collection and hunting, shelter construction, childcare, and so on. This also required an increase in abstract reasoning, an increase in conscious doing accompanied by a decrease in conscious being. Innocent being-in-the-world, which is still the state of the vast majority of other animals, was lost in the process of acquiring language, and man started down a long path of abstract thinking, leading to the knowledge gained through philosophy, the sciences, maths and logic. At the same time, concrete reasoning still went on, in the arts, drama and sports, but already the two strands of human psychology were separated and tended to be emphasised in different individuals.

3. As time passed, the pinnacle achievements of man became associated with conscious doing - abstract thought - rather than the products of contemplative being. At its height, this was reflected in a scientific optimism that all things would eventually be understood through the pursuit of science based in rational thought. The Arts became personal statements by individuals and largely peripheral to the thrust of modernity. God was declared dead (or at least infirm) and religion relegated to one day a week, at best. The separation of our minds from our bodies, our reason from our emotions, was complete. But at some point, perhaps as a result of world wars whose horror was largely identified with technological developments (such as weapons of global destructive capacity), people lost some of their optimism and enthusiasm for the products of conscious doing. The time was ripe for attempts to recombine being and doing, minds and bodies, rationality and feeling, concrete and abstract, thought and action.

4. Computer technology, and associated communication capabilities, were largely a result of efforts to win wars. However, and surprisingly, this most technological and abstract of inventions provided the means for a reintegration of being and doing. Virtual Reality, through its creation of a concrete world from abstract data, opens up the possibility of man recovering his being while not forfeiting technological progress. This is technologically-enhanced, post-literate man, in the situation where computer systems do much of the work of abstraction, and leave him free to contemplate and manipulate the concrete results. That contemplation and manipulation will itself provide conscious scope for new insights and abstractions - or so the educational technologists hope and predict.

To quote Rose (1992, p.93), "Modern technologies - photography, film, video and audiotape, and above all the computer - restructure consciousness and memory even more profoundly [than the introduction of printing], imposing new orders upon our understanding of and actions upon the world". Rose sees this as a step further than printing, which stabilised uncertain observations into "facts". We, however, suggest that while this was true up to some point of history, the development of virtual realities effectively takes us back beyond pre-printing, to pre-literate modes of being.

This is not to say that the world of man will revert to prehistoric times - this is no pre-technological Luddite vision. Rather, man's consciousness is reverting to an earlier state, a state in some ways of innocence, where consciousness is mostly concerned with dealing with apparently-concrete objects (physical doing) rather than mental doing - the hard conscious work of making concrete sense of abstract information - that has been a dominant feature of mental life since language (and, perhaps especially alphabets; see Shlain 1998) first emerged. The impacts of this trend towards presence on what we understand education to be, and thus on how we design and conduct it, are profound.

Presence is a strength in educational settings because the attentional focus produced by a strong experience of being can be expected to activate and motivate the learner. But it is a weakness if attending to the present display inhibits the formation of more general, abstract concepts about the kinds of entities portrayed in the display. In other words, it may be that presence stimulates initial conscious experience in a desirable way, and thus supports relevant perceptual learning, but that it tends to inhibit the kind of conceptual learning that underlies generalisable, abstract knowledge. Such knowledge depends on absence, as well as presence. To achieve timely opportunities for the absence of mind underlying conceptual learning, it is necessary to break the sense of presence at appropriate points in the process.

3 Going Places

Exploration is intrinsic to the presentation of information as virtual worlds, and exploration also implies presence. We cannot explore a place while we are mentally absent from it. Exploration in virtual worlds is a relatively new phenomenon. Although some behavioural studies have considered user navigation in this context, most existing literature discusses navigation in the physical world or hypermedia. Relevant physical-world research has been done in the fields of cognitive anthropology, cognitive psychology, and urban design.

Waterworth (1996) proposed a model for a public information space, called Information Islands, based on hierarchical geographical and urban metaphors. Initial trials were favourable, and this concept is currently being further developed through software prototyping and user experiments, some of which are described below. Related work by Dieberger has demonstrated the navigational and semantic value of urban metaphors for information spaces (Dieberger, 1995). His research has considered primarily text-based interfaces to such spaces, and further work is required to generalise these findings to graphical environments. Recent research by Darken and Sibert has investigated navigation in large-scale virtual worlds (Darken and Sibert, 1996). Their work demonstrated the applicability of some real-world design principles to virtual worlds, as well as the necessity for global structure to support exploration. This research utilised large-scale geographical metaphors; related research in more densely structured virtual environments would extend these results.

Current research by our colleague David Modjeska seeks to explore design tradeoffs between text-based and objectbased representations of information structure that supports user exploration in virtual worlds. Specifically, how do people perceive and learn such representations in VR? What are the implications for navigation and browsing, as well as virtual world design in general? In general, is VR an effective user interface for these information tasks? An ongoing series of experiments is investigating these issues. To support this research, a set of three virtual worlds was developed for desktop virtual reality. This work was carried out at the Department of Informatics at Umeå University in Sweden. The virtual worlds reflected key points along a continuum of design tradeoffs between text-based and object-based representations. Developed designs range from a virtual city landscape to a text-based hierarchy browser. All worlds present the same data, which is a filtered subset of a large, hierarchical index of WWW sites. This data has general interest, rich details, and computational tractability. 1500 information items were included so as to strike a balance between offering a large world for user exploration and completing development in a reasonable time.

Three parallel virtual worlds were developed, with the same input data and maximally isomorphic design features. The first world is an urbanised landscape, with strong spatial and weak textual features (Figure 6). Like a noon landscape, this world has strong colour and lighting cues. The second world is similar, but with weaker spatial and stronger textual elements (Figure 7). Like a dusk landscape, this world offers weak colour and lighting cues. The third world (Figure 8) has weak spatial and strong textual features. Like a city at night, this world has abstract space, with relative but not absolute position. For reference, let us call these four environments the Day World, Dusk World, Night World, and Flat World.



Figure 6: - Day World

The first world was developed according to the Information Islands model (Waterworth, 1996). Data are displayed at six levels of hierarchy using geographical and urban metaphors of island, country, city, district, building, and floor. Additional levels of hierarchy can be added, if needed, through metaphorical elements such as archipelagos and rooms. Design elements suggested by Lynch (1960) were used to articulate space to improve mental mapping, as was global structure as recommended by Darken and Sibert. A number of visual strategies were used to present structure in this way. For example, every parent node in the information structure is represented by a distinctive visual landmark, such as a geometrical object; most boundaries between structure groups are signalled by borders, such as virtual mountains, rivers, or seas. Varied forms and colour groupings help to distinguish object-based entities and regions. Directional lighting enhances the user's sense of overall orientation.

Shifting the balance between object-based and text-based representations, the dusk (hybrid) world is a transitional point between a virtual cityscape and a text-based hierarchy browser. The world is intended to occupy a middle ground, in which the user's mode of perception may shift between text-based and object-based, in the style of an optical illusion. The topography matches that of the object-based world, but object-based features have been reduced in visual intensity. They are represented in unsaturated colours with partial transparency. The labels, by contrast, are bold and colourful. From a design point of view, the Dusk World is intended to offer the best of the Day and Night Worlds.



Figure 7: - Dusk World



Figure 8: - Night World

The study's third world, a text-based one, provides an environment in which the user browses an information hierarchy without object-based features. Semantic cues should predominate over object-based ones. Visually, the world resembles Earl Rennison's "Galaxy of News" (Rennison, 1994) and Apple's "Hot Sauce" prototype. Like Rennison's, the world displays pieces of text in a black 3D space. The user moves between levels of hierarchy and individual nodes by navigating in abstract space. As in Apple's world, the spatial relationships are fixed, but without cues for absolute location or distance. The text labels match those of the Dusk World precisely in location, orientation, size, font, and colour (and the labels of the Day World in every way except colour saturation). Sibling data nodes share the same colour, while adjacent structural groups are assigned contrasting colours for clarity. Directional lighting is not used in this world, in order to make the space more abstract.

The most striking observation we can make about the use of these worlds is the lack of relationship between sense of presence and performance. Performance was not significantly affected by either world design or exposure order, whereas sense of presence was directly affected by world design. Ease of use and overall preference correlated significantly with both presence and performance, yet neither of these latter variables correlated with each other. It appears that the three proposed VR designs are equivalent in usability. However, subjects had a stronger sense of engagement with the Day World, which could have potential benefits in concentration, motivation, and long-term learning.

One criticism of all the environments created is that, although the student is active in exploring the terrain, he cannot change anything. The power of presence is magnified when interaction affects not only position in a virtual space, but the nature of the space itself.

4 Learning by Changing the World

A general model of how humans gain knowledge (Figure 9) is that novices start to become learners by getting declarative knowledge. This means that the novice gets facts about the world to start to create a mental model of the phenomena to be understood. For example, to learn a new language we might start by learning separate words that later can be put together into simple sentences. The beginning learner gains small chunks of knowledge that are stored in memory, although the knowledge consists of small separated fragmented bits of knowledge that are not necessarily linked together. When the learner gains more knowledge about the object in question the chunks of knowledge get bigger and more coherent. At this stage the learner is capable of learning more procedural knowledge - that is knowledge about how to do things, or so called functional knowledge.



Figure 9: A general model of how humans gain knowledge

Whether the learner becomes an expert depends on the task, the frequency of using the newly learned skill and on the learner's motivation and interest. The expert has reached an autonomous stage where the task/skill performs automatically. Tasks thus become easier and faster to perform with less errors. This behaviour is what Rasmussen (1986) calls skill-based behaviour and can be performed by the body semi-consciously or even unconsciously. The learner on the other hand uses active cognitive skill, which mean conscious doing with the mind, which Rasmussen (1986) calls knowledge-based behaviour and which is slower and more prone to errors.

The learning process can be split into three different phases (Figure 10), where the first phase is picking up information that is processed and coded into the memory. The second phase, retention, starts from coding and reaches until the information is activated again; or to put it differently it is when the information is stored passively in memory. Activating the information occurs when the stored knowledge is retrieved. This is done by remembering or by recognising. Norman (1988) terms these "knowledge in the head" (remembering) and "knowledge in the world" (recognising). It is easier to recognise than to remember.



A shortcoming in any one of the phases can result in an inability to remember relevant information. How does this model of learning apply to virtual environments?

4.1 Acting in the Virtual World

The first step towards the direct engagement that conveys a sense of presence in a virtual world was direct manipulation of objects on the computer screen. Although not obvious at the time, we can see now that a change from an abstract, language-based way of interacting with computers to one where computer entities (such as files) and processes (such as delete) are shown as directly manipulable objects (file icons and trash cans) was the first step in a profound shift towards concretising the abstract by computer, rather than by conscious effort. And when people are given things they manipulate, they do things with them, change the way they are organised, and find they have a new undertanding. But as we have seen above, for this to apply to general learning, they must first have acquired at least the basic concepts during the coding phase (Figure 10).

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"). 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. Recent evidence suggests that conceptual processing is controlled through a degree of anchoring in planned sensori-motor processing (see, for example, Hendriks-Jansen, 1996). Findings such as these reflect growing dissatisfaction with cognitivist assumptions within Cognitive Science (see Johnson, 1987, for example).

The idea of two streams of mental activity is implicit in much existing work following the largely cognitivist Human-Computer Interaction (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, 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. In either case, the direct manipulation paradigm strongly primes the user to perceive only a limited set of ways of approaching a problem, as defined by what seems possible at the interface.

In an experiment by Kieras and Bovair (1984) one group of subjects was given a pre-formed conceptual model of a system, while another group did not get any conceptual model, so they had to learn the system by manipulating it mechanically. Both groups received the same procedure training. The conceptual model described how the system was causally connected. Subjects in the group with the conceptual model learned to manipulate the system faster and more easily than the other group.

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 and Waterworth, 1998) 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.

Evidence for the importance of a more "physical" conception of mind than is allowed by cognitivism comes from, amongst other places, 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, 1997). 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 - often forever. This "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.

4.2 Learning and 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 concrete processing is largely unconscious it is not, of course, entirely so. We are often 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 from 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 (reflective cognition).

Changes in this balance between abstract, reflective cognition and concrete reasoning affect the nature of our experience of the world around us. For example, when our conscious processing load is heavy (during difficult abstract reasoning), our experience of duration is short - "time passes quickly" (Waterworth, 1983). We pay little attention to our bodies or the world around us, we are "absent minded" and do not feel present. And when our conscious processing load is light, duration seems long - "time passes slowly" and we are highly present in the current situation - we frequently sample what is going on around us.

Typically, we trade off stimulus sampling with conscious processing as we switch between the abstract and the concrete. Figure 11 illustrates this graphically. We can see the non-processing stimulus-sampling zones of the loops as windows on present reality, and the size of the processing zones as an indication of the degree of abstract processing.



Figure 11: - Location of Consciousness: switching between abstract thought and concrete reasoning

Anchoring the abstract in the concrete reduces the demands on our limited attention span, our consciousness. Losing the abstract does not (usually) result in our losing track of planned physical actions (the concrete). But if we lose track of our mental stream of abstract thought, re-enactment of the most recent physical sequence of actions enables us to recover the mental sequence. In other words, doing things can help us think and learn, provided we get the balance right.

We have recently suggested that different levels of abstraction in virtual environments impose differing burdens on conscious information processing. In particular, cognition of structure from spatial features may be handled largely unconsciously, while cognition of the structure conveyed by textual features is probably handled consciously. These differences should result in corresponding changes in conscious capacity available for other activities during exploration. If so, such differences should be reflected in a user's completion of informational tasks during exploration, as well as in his/her recall of environmental structure afterwards. Assessing performance and recall levels would thus measure the cognitive difficulty of exploring virtual environments of differing abstraction.

In order to design for learning one adequate characteristic to stress is to try to engage the learner. One way of doing this is to help the learner develop simple procedural knowledge that he can use in a practical situation but at the same time be able to learn declarative knowledge. The learner has to see a fast result of the learning process and that he gains useful knowledge that is useful in the reality. This implies an intertwining of procedural and declarative knowledge during the learning process. Most commentators would agree that this mix is not fixed during the whole learning process. In the beginning procedural knowledge should be stressed in order to engage the learner, while later in the learning process more declarative knowledge should be stressed - which amounts to satisfying the learner with more useful knowledge.

5 Conclusions: Presence and Absence

It may never be possible to define presence, but it is possible to identify dimensions of variation affecting the overall sense of presence in particular VRs. In Waterworth and Chignell (1991), a model of human information exploration was characterised along three axes. This model has stood the test of time, with a few slight changes of terminology. We now think of information exploration activities as being located in a space defined by the dimensions of structural responsibility, interaction loop speed, and specificity of user requirements. These dimensions reflect how individuals deal with the information that is the product of exploration, and in this, a consideration of consciousness and the emphasis of conscious activities is central. This is not surprising of course, since mind can be viewed as an exploring and storing system with the function of delivering relevant information to consciousness.

Figure 12a represents an attempt to identify the main dimensions of mind. The high-low focus axis refers to the extent to which our attention is directed to fine-grain detail or the broad stroke features of a situation (see Gelernter, 1994). This may be our sampling of the environment or our conscious processing of previously-sampled information. When our focus is high we might be trying to see a particular distant object on the horizon (a particular kind of animal, for example), or to solve a problem with mental arithmetic. When our focus is low we might be daydreaming or admiring a landscape. The conscious-unconscious axis refers to how conscious we are. This is often correlated with level of "wakefulness" although we may actually be largely unconscious while awake and highly conscious when asleep (as in vivid dreaming). The frequent-infrequent sampling axis refers to how frequently the individual samples from the stimuli received by the senses (as illustrated in Figure 11). Rapid sampling will tend to occur when the conscious processing load is light, or we have for some reason to be highly vigilant to what is happening around us, and will be accompanied by the experience of time passing relatively slowly. Infrequent sampling typifies when we are fully engaged with conscious processing (planning a future public speech, for example) and take little notice of our physical circumstances (internal or external).



Figure12b: Dimensions of Information Exploration

Figure 12b illustrates the information exploration model of Waterworth and Chignell (1991). The high-low specificity axis refers to how clear the explorer is about what is being sought - a specific information query versus browsing around. This corresponds to the high-low focus dimension of mind. The system-user responsibility axis refers to whether the system or the user does the searching (and must therefore be aware of the information terrain). This corresponds to the conscious-unconscious dimension of mind. The rapid-slow interaction loop refers to the style of interacting with the information system. Rapid loops typify interactions based around direct manipulation and the manual following of links by the user. Slow loops arise when the user adopts a conversational style, perhaps describing a complete information request to an automated agent that will return with a response some time later. This corresponds to the frequent-infrequent sampling dimension of mind.

The physical landscape contains and supports cognition. In the same way, a virtual landscape will support conscious processing; its nature will help determine the location of consciousness, as well as its content. When people are given things, they do things with them. The educational question is whether they learn anything useful in the process.

Virtual environments are actually ideal aids to engaging the user in the learning process, but they should not only be for exploring (conscious being), but also to encourage abstract thinking (conscious doing). This implies directing the

development of humans from conscious beings via doers to changers. Changers are people who integrate being (the concrete) with doing (the abstract) in order to gain better understanding.

A sense of presence during exploration will be experienced most strongly when the learner has structural responsibility, the opportunity for fast interaction with representations of information, and when he is not very clear what he is seeking. This typifies highly conscious, unconsolidated knowledge, and is only part of the cyclical pattern inherent in long-term conceptual learning. To complete that process, breaks in the sense of presence are also needed. Education is not only a process of information directed exploration, it also requires absent-minded wandering.

7 Acknowledgements

We are very grateful to David Modjeska for his contributions to the work reported here.

8 References

Allwood C.M. (1989). Computer usage by novices. In: Kent I.A., Williams J.G.: *Encyclopedia of Microcomputers*. Marcel Dekker IMC, New York.

Anderson J. R.(1983). The architecture of Cognition, Cambridge, Mass.: Harvard University Press.

Darken, R.P. and J. L. Sibert (1996). Wayfinding Strategies and Behaviors in Large Virtual Worlds. In *Human* Factors in Computing Systems: Proceedings of the CHI '96 Conference. New York: ACM, 1996.

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.

Dieberger, A. (1995). Providing Object-based Navigation for the World Wide Web. In Andrew U. Frank and Werner Kuhn, eds., *Object-based Information Theory: A Theoretical Basis for GIS: International Conference COSIT '95*. Berlin: Springer, 1995.

Dreyfus H.L., Dreyfus S.E. (1986). *Mind over machine. The Power of Human Intuition and Expertise in the Era of The Computer*, Cambridge, England: Basil Blackwell Ltd.

Gelernter D. (1994). The Muse in the Machine - Computerizing the Poetry of Human Thought. The Free Press, New York

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.

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

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

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.

Kieras Bovair (1986). The role of a Mental Model in Learning to Operate a Device. Cognitive Science 8

Laird J.E., Newell A., Rosenbloom P.S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence* 33(1):1-64

Lund, A. and Waterworth, J. A. (1998). Experiential Design: Reflecting embodiment at the interface. Presented at: *Computation for Metaphors, Analogy and Agents: An International Workshop,* University of Aizu, Japan, April 1998.

Lynch, K. (1960) Image of the City, The. The MIT Press, Cambridge MA.
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.

Newell A., Rosenbloom P.S., Laird J.E. (1989); "Symbolic Architecture for cognition",

In: Posner M.I. (Ed); "Foundations of cognitive science", Bradford Book The MIT Press Cambridge Norman D (1988). *The Psychology of Everyday things*. New York: Basic Books Norman D (1993): *Things that make us Smart*. Addison-Wesley Publishing Company Norman, D. (1986). Cognitive Engineering. In Norman, D and Draper, S., *User Centered System Design*. Hillsdale: Lawrence Erlbaum Associates.

Rasmussen J. (1986) Information Processing and Human-Machine Interaction - an approach to cognitive engineering. Series Volume 12, North-Holland.

Rennison, E. (1994), "Galaxy of News: An Approach to Visualizing and Understanding Expansive News

Landscapes." In UIST '94, ACM Symposium on User Interface Software and Technology. New York: ACM Press, 1994. Rose, S. (1992). The Making of Memory. London: Bantam Press.

Shlain, L (1998). *The Alphabet versus the Goddess: The Conflict Between Word and Image*. New York: Viking. Shneiderman, B (1988). *Designing the User Interface*. Reading, Massachusetts: Addison-Wesley.

Tulving E., Pearlstone Z (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 5, 381-391

Waterworth, J. A. (1983). *The Influence of Variations in Cognitive Processing on the Perception of Time*. PhD thesis, University of Hertfordshire, UK. Available through the British Lending Library (Accession no. D50267/84).

Waterworth, J. A. (1996), A pattern of islands: exploring public information space in a private vehicle. In P. Brusilovsky, P. Kommers and N. Streitz (eds.), *Multimedia, Hypermedia and Virtual Reality: Models, Systems, and Applications*. Lecture Notes in Computer Science 1077. Berlin: Springer-Verlag, 1996.

Waterworth, J. A. (1997). Personal Spaces: 3D Spatial Worlds for Information Exploration, Organisation and Communication. In R. Earnshaw and J. Vince (eds.): *The Internet in 3D: Information, Images, and Interaction*. New York: Academic Press, pp. 97-118.

Waterworth, J. A. and Chignell, M H (1991). A Model of Information Exploration. *Hypermedia*, *3* (1) 35-58. Yates, F. A. (1984). *The Art of Memory*. London: Routledge and Kegan Paul. First published in 1966.