

Presence and Absence in Education VR: The Role of Perceptual Seduction in Conceptual Learning

John A. Waterworth and Eva L. Waterworth
jwwaterworth@informatik.umu.se

Dept of Informatics
Umeå University
SE-901 87 Umeå
Sweden

Abstract

Recently developed information technology, such as interactive multimedia and virtual reality, has a profound impact the way people deal with information. 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 this paper, we discuss the respective roles of presence and also of absence in relation to virtual worlds designed for learning, and to existing theories of learning. We characterise absence as a psychological focus on conceptual processing, and presence as a psychological focus on direct perceptual processing (of things which are present in the current environment). We suggest that virtual worlds tend to elicit a sense of presence because they present information as directly perceptible objects and spaces, rather than as linguistic specifications requiring conceptual processing if they are to be understood.

To investigate this idea, a series of experiments was conducted in which the degree to which information was specified as objects or text was varied. An account of our findings is followed by a description of our model of the role of presence and absence in virtual environments designed to encourage conceptual learning. We conclude by suggesting that effective conceptual learning in virtual worlds depends on designs that support both perceptual and conceptual processing – worlds that provide both the necessary perceptual cues to elicit a sense of presence, and that contain a sufficient element of surprise to break the spell periodically and so foster conceptual learning. The resulting changes in how students deal with relevant knowledge promise to make classroom learning less exclusively conceptual, but also less fragile, and more like learning in the world outside.

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 typical of interacting with new technologies and the learning process. After that, Section 3 introduces and expands on the idea of information access in learning situations as exploration of an information landscape. Section 4 emphasises the importance of students being active in achieving tasks in the virtual realities produced by new technology. 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 present our model of the respective roles of presence and absence in learning in general, and education in particular, before drawing our final conclusions in the last section. 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 into working memory and long-term memory. Act* divides long-term memory into two parts, declarative 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 et al., 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.




Figure 1 - Different types of knowledge representation in memory.

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.

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. Pre-learners and experts can be said to be in a state of conscious being; they are present with the phenomena with which they are

dealing. Novices, 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 the presence-absence distinction).

Three levels of control of human actions




Figure 3 - Three levels of control of human actions (based on 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. 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 pre-learners. This is because neither pre-learners nor real experts are conscious of the structured knowledge underlying the topic with which they are dealing. In the case of 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 formal education, which still emphasises theoretical learning and abstract thinking. 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 qualified tailor. We will suggest later (Section 5) that the different types of knowledge that these two approaches aim to instil in learners require a different mix of mental presence and absence, and that this can inform the process of designing appropriate VR learning environments.

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

Figure 4 - The two extremes of a spectrum of learning situations

1.3 A Note on Terms

The distinction between concrete and abstract processing is crucial to an understanding of presence and absence of mind, and of how VR can support education that is not merely training of the body. Various terms are used to make this rather elementary distinction in the literature. We consider the following terms within two sets to be synonymous for the purposes of this paper. Firstly: abstract, conceptual, and reflective processing (or cognition) all refer to what we also call absence of mind; that is dealing with concepts which stand for things which may or may not be currently present (and which has also been called "knowledge in the head"). Secondly: concrete, perceptual and experiential processing (or cognition) all refer to what we call presence, that is dealing with things that currently exist in the external or internal environment (and which has also been called "knowledge in the world"), including sensori-motor responses.

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. Conceptual learning is all about the generalised, the abstract, that which is physically absent. Perceptual learning is all about dealing with the present here and now in concrete ways. 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. But even conceptual skills become

largely automatised with time. The expert speaker of a foreign language, for example, does not need to think about the grammatical rules underlying competent use of the language.

New information 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 (more details of this "history of being" can be found in Waterworth 1999a). One of the areas where this is most actively resisted is in formal education. These technological developments explain the recent spate of interest in the sense of presence - in what factors generate it, and in its effects.


2.1 Views of presence

Presence is the hallmark of virtual reality. VR designers typically aim to create a convincing, engaging environment in which users feel present. When this feeling of "being there", immersed in a virtual world, is strong users do not seem to have to conceptualise about the world to make sense of what is portrayed. One interpretation of this is that such virtual environments do not require mental modelling (knowledge in the head) to be make sense. Rather, they contain "knowledge in the world" which is amenable to direct perceptual processing. The gold standard of VR can then be seen as passing the "VR for Animals" test (Waterworth, 1996a) - could the VR convince a non-human mammal (since such animals exceed human perceptual skills, but lack most of our conceptual capabilities)? And VR is not limited to attempts to duplicate the real world.

There is certainly a growing interest in the notion of presence; in entertainment, media studies, professional training, informatics, and of course education. While TV broadcasters and professional trainers are happy to maximise presence, educationalists worry that to experience something vividly does not always result in useful learning. This interest in presence in VR has led researchers to focus on issues such as the ways in which presence can be assessed, and on models which attempt to explain how the factors affecting the evaluated level of presence in specific VR implementations. Assessment of the sense of presence is an almost universal feature of recent approaches to evaluating virtual environments. Most commonly, this is done by asking VR users how present they felt - usually by means of some subjective scaling procedure.

Waterworth (1999b) presents a spectrum of VR application types and their evaluation, as shown in modified form in Figure 5. The spectrum ranges from the most abstract (abstract-external) to the most concrete (concrete-internal). At the far right, the concrete-internal perspective emphasises concrete, bodily experiences; changes due to the VR that can be measured physiologically. Here, presence is very much the focus, and we could see physiological change as a candidate measure of presence. An example is in interactive art works that engage the immersant in a vivid, present reality. Another is in the way VR can be successfully used in the treatment of anxiety disorders, such as phobias, which require a re-programming of emotional bodily responses to certain environmental stimuli. At the other end of the spectrum, the abstract-external perspective emphasises expert assessments of outcomes, and reflective processing of information. Here, a sense of presence is treated as incidental. Educational applications of VR typically fall between these two extremes.

While it is generally accepted that feeling present is motivating for students, there is also disquiet that vivid experiences are not sufficient for conceptual learning; to experience is not to understand. So while presence is obviously necessary for training concrete skills, as in a driving or flight simulator (concrete-external), there is some debate about the role of presence in conceptual understanding (abstract-internal). Such understanding involves the internalisation of an abstract body of knowledge, whereas training systems develop appropriate concrete responses of the body itself.



Spectrum of VE applications and Evaluation Types

Figure 5 - A spectrum of VR application and evaluation types

Models of the role of presence are in general somewhat underdeveloped. Slater et al. (1994) suggest that presence is not the same as immersion, but that high immersion tends to yield high judgements of presence. Immersion depends on the technology used in the VR - a head-mounted display is more immersive than desktop VR. They suggest that presence, rather, is a state of consciousness. While this seems plausible, it is not very explanatory: the question of when does more immersion yields greater presence, and when it does not, is not answered. Other models, while more predictive, are not always plausible. Bystrom et al. (1999) present the Immersion, Presence and Performance (IPP) model. Roughly speaking, this suggests that immersion produces sensory fidelity, which directs the allocation of attention in such a way that subjective disbelief is suspended, yielding high presence and, as a result of that, better task performance. While this model has the advantage of explicitly relating presence to both level of immersion and the allocation of attention it seems to us to be obviously wrong. The implication that there is no performance without a sense of presence is obviously false. In our own experiments (Section 3) we found that performance was better in a hypermedia environment in which students felt no presence than in an equivalent virtual world where they felt highly present. Clearly, people performed tasks before the advent of virtual worlds! Their model also suggests that there will be no sense of presence without sensory fidelity, but this is also not borne out by the literature.

Witmer and Singer (1994) attempted to separate out the technological versus attentional aspects of the sense of presence in a slightly different way. They posit that spatial presence is a direct product of the level of technological immersion, whereas

attentional presence arises from the extent to which users become involved in what the virtual environment portrays. Presence thus depends to some extent on task demands.

The root of the problem with existing models of presence is perhaps a confusion between presence and suspension of disbelief. Suspension of disbelief is the result of conceptual processing which then leads to a secondary sense of involvement - as when we read a gripping novel in which we become engrossed. On the other hand, we see presence as that which arises in situations where no belief suspension is needed - because the display is immediately perceptually engaging (though not necessarily faithful to normal sensory perception). These two tend to be confused in VR studies because of the dominant method of presence assessment by subjective ratings. Sometimes subjective judgements are in conflict with physiological and physical responses. A person may say that a VR display is very unrealistic and does not fool them, but may still show physical responses (such as ducking out of the way of an approaching virtual object) and physiological correlates (decreased skin resistance, quickened pulse). The reader of a novel may become deeply engrossed in the lives of the characters and the action that is described, but they are unlikely to move their bodies unconsciously to avoid a hazard that is only described in text.

Schubert et al. (1999), see presence as "embodied cognition". They suggest that presence emerges when bodily responses are conceived as possible in a VR. This is compatible with the approach of Stappers et al. (1999) who adopt an ecological approach to human behaviour, one where they look for objective performance correlates of subjective reports. For example, whilst low sensory fidelity (more abstract images) was found to result in low ratings of presence, people often behaved bodily as if they were really present in the virtual situation. They conclude that subjective reports need to be supported by performance measures of presence - such as body movements. Similarly, Zahorik and Jenison (1998) see perception-action coupling as necessary to a sense of presence, or what they call (after Heidegger) "being-in-the-world".

Despite reservations about existing models of presence, it seems necessary to consider experienced presence when trying to understand how VR environments might help develop conceptual learning in students. For example, Salzman et al. (1999) present a model that purports to explain this. But by failing to consider different aspects of the conscious experience of the learner, their framework comprises merely an interconnected web of situational characteristics. We consider presence, which we see as a reflection of the balance of activity in the consciousness, to be a key element in understanding the role of VR in learning in general and education in particular. In the next subsection, we outline our own model of presence and absence - the two sides of the same coin of conscious attention.

2.2 Presence and the Location of Consciousness in VR

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 concrete (or perceptual) and the abstract (or conceptual). In other words, our capacity for conscious attention is shared between presence (perception of the environment) and absence (processing of concepts).

One way of characterising the dichotomy between the two streams of mental life is as a question of where consciousness is located. Although perceptual processing is often largely unconscious it is not, of course, entirely so. We are often aware of what we are feeling and doing in the sensori-motor world of the concrete here and now. And while conceptual processing is usually 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 (concrete, sensori-motor processing) and carrying out conscious mental work (reflective or abstract conceptual processing).

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 share conscious attention between the abstract and the concrete. Figure 6 illustrates two situations with varying balance between concrete and abstract processing.




Figure 6 - The location of consciousness: two examples

At the right in Figure 6, the emphasis is more on stimulus sampling from the environment; to the left the emphasis is on conceptual processing. These stimulus-sampling zones of the loops can be seen as windows on present concrete reality, and the size of the processing zones as an indication of the degree of abstract processing. When the emphasis is on stimulus sampling, the experiencer is more present in the current environment and time will seem to pass slowly. When the emphasis is on abstract conceptual processing, the experiencer is more absent and time seems to pass relatively quickly.

To make the two examples in Figure 6 more understandable, imagine that you are in a hospital room visiting a patient, a friend who is paralysed but has a wonderful memory. There is a telephone in the hall outside your friend's room, and he needs you to make some calls for him. Neither one of you has a pen or paper, so he must tell you the number to dial, and you must keep it consciously in mind while you walk out to the telephone and dial. In the first example, to the right in Figure 6, the number is long-distance and therefore long, say ten digits (e.g. 091-847-3062). You repeat this number

to yourself as you walk to the hall phone. In this case, your conscious capacity for attending is largely taken up with abstract processing (articulating the numbers in consciousness). Because of this, you have little capacity left for sampling concrete details from the environment around you, and time passes quickly for you. You do not notice out the environment around you, rather like the stereotypical "absent-minded professor" because your consciousness is mostly occupied with things other than the present environment.

In the second example, to the left in Figure 6, the number is local and short, say five digits (e.g. 78638). Now the conceptual processing task is relatively light and easy, and you have conscious attention to spare as you walk to the phone. This time you notice more from the present environment, and time passes relatively slowly for you. You observe that you are hungry and that there is a sandwich machine at the end of the hall, aspects of the environment that had escaped your attention on your first trip down the hall.

The little story above illustrates our model of presence and absence. Presence arises when we mostly attend to the currently present environment within and around the body. The capacity we have for such attention depends on the amount of conceptual processing the situation demands. As we process more in an abstract way, we can consciously sample fewer concrete aspects of the present situation, and so our sense of presence diminishes; we become absent. Figure 7 illustrates this trade off in more general terms, viewing the mind as a two-roomed apartment lit by a single light source. The rooms represent concrete and abstract processing, respectively. The light is consciousness, and areas in darkness are unconscious.




Figure 7 - The mind as a two-roomed apartment

If the light is angled to the left, concrete processing of the present situation is emphasised, and there is a strong sense of presence. At the extreme, an observer in this state of mind has no capacity left for abstract thought. He is having a vivid experience that seems to take a long time to unfold, though he will have little recallable memory about it. If the light is angled more to the right, however, abstract processing of concepts is emphasised, and there is little or no sense of presence. The individual has little awareness of what is going on around him. Time passes quickly, and his

experience is only of the concepts with which he is dealing and of which he will have a good recallable memory.

VR can elicit presence, just as physical reality can. And perhaps the most obvious way of using VR in the educational process is by creating virtual landscapes of relevant information which students are encouraged to explore in various ways (Waterworth and Waterworth, 1999).

3 Education as Exploration

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, and is the most common way of allowing VR users to interact with sets of information. 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 (1996b) 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.

Recent research by our colleague David Modjeska seeks to explore design tradeoffs between text-based and object-based 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 the Yahoo! 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 8 - Day World). 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 9 - Dusk World). Like a dusk

landscape, this world offers weak colour and lighting cues. The third world (Figure 10 – Night World) has weak spatial and strong textual features. Like a city at night, this world has abstract space, with relative but not absolute position.



Figure 8 - Day World

The first world was developed according to the Information Islands model (Waterworth, 1996b). 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 (1996). 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.

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.

We can conclude from these experiments that different levels of abstraction in virtual environments impose differing burdens on conscious information processing. In particular, cognition of structure from spatial features involves largely unconscious concrete processing, while cognition of the structure conveyed by textual features involves more conceptual processing and takes up conscious capacity that might otherwise be available for other conceptual work. 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. Estimates of duration are also useful indicators of the level of conscious conceptual work carried out during an interval (Waterworth, 1983).

4 Learning by Changing the World

One criticism of all the environments described in the last section is that, although the student is active in exploring the terrain, he cannot change anything. This is appropriate, however, for novices (note the two first stages of the model of Dreyfus and Dreyfus, 1986). Following standard constructivist views of learning, we would expect that the power of VR to encourage conceptual learning would be magnified when interaction in the VR affects not only apparent position in a virtual space, but the nature of the space itself (the third stage of Dreyfus and Dreyfus, 1986). In following this line of argument, we look first at how learning tends to occur in the real world.

4.1 Learning in the Real World

In order to design for learning one necessary 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 novice, while later in the learning process more declarative knowledge should be stressed - which amounts to satisfying the more expert learner with more useful knowledge.

A general model of how humans gain knowledge (Figure 11) is that pre-learners start to become novice learners by getting some initial 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 novice 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 11 - 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 novice on the other hand uses conscious conceptual skills, which Rasmussen (1986) calls knowledge-based behaviour and which is slower and more prone to errors.




Figure 12 - Learning: three phases

The learning process can be split into three different phases (Figure 12), 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 do these models of learning apply to virtual environments?

Results from a few studies of people who earn their living dealing with abstract concepts, so-called "knowledge workers", provide a starting point for designing conceptual learning in VR.. It has been claimed that these paradigmatically conceptual 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, 1997a). 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 this knowledge has been "understood" enough to be filed away, it is ready to be forgotten - to become inaccessible to consciousness. 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 conceptual in the sensori-motor, the bodily side of perception.

Anchoring the abstract in the concrete reduces the demands on our limited attention span, our consciousness, as discussed in Section 2. Losing the abstract does not (usually) result in our losing track of planned physical actions (the concrete). And 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. The conclusion we draw about real-life knowledge workers is that they are more concrete, and less abstract, than we might have supposed. Because of the inherently concrete nature of VR (which is what makes it seem like a "reality"), when we design for learning in VR we should aim to capture the ways in which learning in the real world (rather than the classroom) takes place.

4.2 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 understanding. 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 12).

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.

Golightly et al. (1996) cite studies showing that, sometimes, direct manipulation results in poorer problem solving than an equivalent command-based interface, 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). Perhaps, as the authors suggest, the direct manipulation paradigm primes

the user to perceive only a limited set of ways of approaching a problem, as defined by what seems possible at the interface.

These results confirmed the earlier study by Schär (1996) who relates his findings to a more detailed consideration of styles of learning. Schär suggests that direct manipulation tends to foster implicit learning, whereas a more conversational interface produces a bias towards explicit learning. Using Tversky's (1980) model, implicit learning is considered to be unconscious, exploratory and creative, whereas explicit learning is conscious, as well as more "rational" and focused. In our terms, explicit learning is a function of abstract processing (which gives mental absence from the interface), whereas implicit learning is the product of more concrete processing (which also results in feeling present at the interface). Schär (1996) found that a direct manipulation interface induced implicit learning, whereas a conversational interface induced explicit learning.

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.

The general pattern, then, of introducing technology which enhances virtual realisation, is to free conscious capacity that would otherwise be needed for mental visualisation in the imagination. Professionals can then function on the basis of the unconscious, automatic or automatised skills that underlie pattern recognition (Nygren et al., 1992) and smooth sensori-motor performance. Because it is possible externally to represent and experience information in new ways - rather than by using the limited capacity of human consciousness for internal representation - this will also open up new possibilities and may stimulate creativity. However, there are also dangers inherent in these changes. Firstly, skills are likely to be lost through lack of use, and will thus no longer be available should the technology fail or be unavailable. Secondly, the results obtained with the new techniques may not always be an improvement on the old. This is because users may be drawn by the technology into using easier but less effective strategies for problem solving.

The physical landscape contains and supports cognition. In the same way, a virtual landscape can 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.

5 Presence and Absence in Education

"Unpredictability is often said to be the essence of creativity. But unpredictability is not enough. At the heart of creativity lie constraints: the very opposite of unpredictability. Constraints and unpredictability, familiarity and surprise, are somehow combined in original thinking."

Margaret Boden (1995)

Virtual environments are ideal aids to engaging the user in the learning process, but they should also encourage abstract thinking to produce real learning. This implies that learners are not only explorers of places or easy solutions, but changers. Changers are people who integrate the concrete world of the present with the abstract possibilities of the absent to gain a better understanding of the object of study.

5.1 Presence and Information

We have presented a simple model of the sense of presence in Section 2. From this 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 13 a: Dimensions of mind, b: Dimensions of information exploration

Figure 13a 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 Figures 6 and 7). 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).

Figure 13b illustrates the information exploration model (slightly modified) 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.

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, biased towards learning the use of knowledge in the (virtual) world, and is only part of the cyclical pattern inherent in long-term conceptual learning. To complete that process and develop new knowledge in the head, breaks in the sense of presence are also needed.

5.2 Presence and absence in educational VR

Some recent VR-like computer games combine presence and absence, exploration and problem solving. But the conceptual work only helps deal with the virtual world itself (puzzling out how to get to the conclusion of the game successfully) and so does not build up generalisable conceptual knowledge.

Today's education stresses abstract and logical thinking (which in our terms is to be absent) and preferable using a high (i.e. large) attention span. One problem this produces is that people are not good at applying high attention to relevant concepts over a long period, since this is very tiring and the mind tends to wander onto other things. Most of today's VR applications, on the other hand, stress experiencing and exploring information, emphasising presence, which makes the users feel that they are there within the created reality. This also tends to involve a high attention span, though focused on the information presentation itself rather than on concepts behind the representation. The weakness of this approach is that although it is initially very engaging, the user does not gain much new knowledge, even when viewing the display for a long period. And after a while, boredom sets in, and once again the mind starts to wander. All of this points out the necessity of every now and then breaking the absence in traditional education and of breaking the presence in VR-based learning. The opportunity to change the world that is modelled in a VR provides a focus for breaking presence - since meaningful change requires conceptual work for its planning and execution.

These breaks occur differently depending on where the learner is in terms of skill acquisition stages. A novice has to direct all of his/her attention to the skill being learnt both when it comes to conceptual and to perceptual learning. Furthermore a novice uses knowledge-based behaviour requiring conscious attention. This points to the necessity of moving between being present and being absent and in this way reduce the heavy workload involved in learning a new task. Another advantage of shifting between presence and absence in the learning situation is that it makes the user both more engaged and involved in the learning situation.

An expert on the other hand does not need to use most of the attention span when adding new items to the existing knowledge because he/she has already both declarative and procedural knowledge about the topic. Also, the expert is able to mostly rely on skill-based behaviours which are largely unconscious or to use rule-based behaviour. Neither of these kinds of behaviour require the learner's whole span of attention. This suggests that the existing model of an expert is only expanding, but in order to keep the expert engaged and feeling motivated to learn more without being bored, it is important to keep transferring between a mental state of presence and that of absence.

From the theories and models described in this paper we suggest a learning model that stresses the importance of breaking the sense of presence to carry the learner forward in the learning process. This model also considers where the learner is in terms of skill acquisition stage. Another important factor is to try not to have long continuous periods where the learner has to use a high attention span. We believe that the model could be used to understand the acquisition of conceptual skills as well as perceptual skills, although the learning situation differs between the two. The main components of the model are the difference between absence and presence, the quantity of attention demanded at the moment, and the five skill acquisition stages (figure 14). Another assumption is the use of VR in order to be able to create the breaks between absence and absence and to facilitate the possibility for the learner to change and interact in an almost natural situation without causing considerable damage.




Figure 14 - A model of perceptual and conceptual learning in VR

The model is an aid to understanding, analysis and the design of learning environments. As mentioned before the model applies both to conceptual learning and to perceptual learning, which only differ in how they are performed. The model only depicts use of attention and the situation of the specific object that is learned.

An example of conceptual learning is, for example, when you want to learn a new language (see figure 15a). At first you have to learn a few single words and some simple rules to put the words together in order to create simple sentences. This takes all of the learner's attention and makes him/her absent - using abstract thinking. This is typical knowledge in the two first steps at the skill acquisition stages, novices and advanced beginner. These stages are very abstract and conscious. Then suddenly the learner starts to experiment (change) with their new learned skill and tries to speak the new language more spontaneously, which indicates that he/she has become competent and that he/she is more present, but still most of the attention span has to be used in order to practice the newly learned skill.

As the learner becomes more and more experienced it is possible to use less and less attention in the interaction, but he/she still has to learn more words and rules which means every now and then to break the presence and return to absence. An expert does not need much attention to learn more because by that stage it is possible to learn at the same time as using the skill. This implies that the jumps between presence and absence are more or less unconscious, but that in a situation of need the expert has the possibility to switch to a high attention span of absence or a high attention span of presence depending on the situation.




Figure 15 a) an example of conceptual learning, b) an example of perceptual learning

When learning a perceptual skill on the other hand, as for example to ride a bike, the learner starts with learning the importance of keeping the balance of the bike and trying to practice that skill on the bike. This takes almost all of the learner's attention span and is performed mainly by the body in a state of high presence (Figure 15b). After reaching the stage of preserving balance it is time to learn to steer, brake and control the bike, which leads the learner in to the next stages of acquisition skill, advanced beginner and competent learner. As the learner improves the skill he/she is able to use less attention riding the bike, and that attention could be used for other things. As in conceptual learning an expert does not need much attention to learn more because in that stage it is possible to learn at the same time as using the skill, so this implies that by the time the learner becomes proficient or expert the curves have become the same.

5.3 The Wheel of Creative Learning: Cycling through Presence and Absence

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, as we have suggested above.

Fencott (1999), in similar vein, points to the importance of both perceptual cues and surprises for conceptual learning in educational VR. Cues perceptually reinforce the unconscious expectations one has when acting within a virtual; in other words, they engender and support a sense of presence. Surprises, on the other hand, break those expectation (and the sense of presence) and stimulate reflection. Whitelock and Jelfs (1999) report compatible experiences with a range of virtual learning environments.

What is true for learning is, we suggest, also true for creative thought. Waterworth (1997b) pointed to the way in which new technologies, by allowing information to be experienced in a variety of vivid media and forms, expand the possibilities for creative inspiration. In other words, experiencing the same underlying information in different representations, and through a variety of sensory channels increases the range of concrete perceptions through which information is experienced. These richer perceptions may then lead to more original concepts. For this to happen, as with conceptual learning, we suggest that there should ideally be both the feeling of presence afforded by rich perceptions, and the absence of mind required to examine any resultant insights conceptually. This gives a new way of looking at education, and at learning in general.

Human beings learn throughout life, the new born has to learn a great many sensori-motor skills, children learn to talk, cycle and a lot of knowledge in school, we learn a profession, to be a wife or husband, to look after children. Even old people have to learn to stop working and adjust to the new lifestyle of being old with everything that involves. All of this shows that learning has a very central position for human beings - to be human is to learn. Education can be said to be a controlled way of learning, as for example in today's school system from early age to different kinds of higher education and separate courses for adults. One problem with today's education, as mentioned earlier, is that it stresses abstract and logical thinking, which we call absence, at the cost of concrete representation and reasoning (perception), which we call presence (see for example Waterworth 1999a, Lindh and Waterworth 1999).




Figure 16 - The wheel of creative learning

We believe that when people learn new things and are engaged and motivated to learn they have to oscillate between presence and absence in a similar way as when they are creative. Our view is that all learning includes a creative aspect in the development of useful and understandable knowledge. This implies that learning and creativity have a lot in common, including the movement between presence and absence, as well as knowledge transfers between consciousness and the unconscious. Our model of learning, the education wheel (Figure 16), is very similar to our model of creativity (see Lindh, 1997). The wheel combines the best of traditional education and VR-applications which elicit presence. Furthermore it combines the strengths from both traditional learning and learning by doing (see Figure 4).

5.4 Two Examples of the Learning Wheel in Motion

The wheel of learning (Figure 16) should be viewed as a wheel rolling forward to new experiences. It does not describe the whole learning process, as for example of learning to cycle or learning a new language. It portrays the ongoing process in learning; the wheel rolls around all the time. This takes place both for novices and all the range of learners up to, and including experts. Furthermore it does not matter much whether the learning is conceptual or perceptual. There is no start or ending of the wheel's journey to new knowledge, but if we think of a specific learning task that will determine where we enter the wheel.

5.4.1 Perceptual Learning in Life

As an example of perceptual (sensori-motor) learning, consider a person who is quite good at cycling, but who has never cycled in winter when it is frozen and slippery

on the roads (here in Sweden, at least). We start with the position where our friend is out cycling and suddenly he experiences the bike starting to slip and he is about to fall. This is a concrete experience where his body and emotion perceive this input. Most of his actions are under the control of the unconscious but our bike rider has to frequently sample the environment consciously. He feels that time passes very slowly and suddenly he gets a spark, an idea, how to handle this new situation. This means he moves from presence to absence to check the new idea.

While absent time seems to pass faster since in order to conceptually examine and create a plan of his new idea he has to stop sampling the environment as frequently as he did before. His consciousness takes control of designing and executing the plan, which means that he uses high focus thoughts, where he concretises the abstract mentally. After the new plan to solve the problem is formed, he has to judge and execute the plan, which leads to experiencing a changed situation, and the wheel continues.

It is easy, in principle, to transpose this learning experience into a VR environment, where all the relevant features of the experience are synthesised and learning proceeds in the same way as in real life, but without the potential for physical damage to the individual concerned. However, it is unlikely that the benefits would justify the costs involved in this case. Similarly true-to-life sensori-motor simulations are exemplified by flight simulators and some rehabilitation systems for disabled individuals.

5.4.2 Conceptual learning in VR

Now imagine that a class of intermediate level students in their history lesson is studying the battle of Waterloo. To learn about Napoleon in general and Waterloo in particular they are using a desktop VR application that is build like a role playing game. Every student has his or her own part to play, for example one plays Napoleon, one plays Lord Wellington (who is in command of the English troops), and one plays Blücher (who was head of the Prussian army). We start our example where the students have arrived at Waterloo with their troops and are occupying initial positions. This kind of action demands presence from the student. He frequently samples the environment in order to keep a watch on the current situation and discover any significant moves by the other players in the battle. He is highly vigilant and experiences the environment richly, acting through unconscious control with the body and emotions largely in charge. Suddenly the game stops and every student gets a written question on their screen, specific to their own concerns in the battle. For example, the students may have to specify their plan of immediate action.

This surprising event breaks the sense of presence and forces explicit conception of the situation. The student no longer samples the environment as frequently as before and his moves to high focus thoughts. This means that time seems to pass much quickly and he starts to create a plan to solve the questions so that he can continue the game. When delivering the answers he checks to see if they are likely to be right, which indicates that he judges the situation. If he answers the questions reasonably he is allowed to continue the game, which switches him back into a state of presence. Suitable cues must be provided by the VR to engender such periods of presence, which motivate and engage the student. Surprising breaks in the virtual action force periods of explicit conceptual learning, rewarded with presence during another episode of more direct experience.

Cycling through presence and absence is more powerful than either in isolation. Pure presence would be engaging, at least in the relatively short term, but would not lead to much explicit conceptual learning, although some implicit perceptual knowledge might adhere to the student. Pure absence - the same facts conveyed through a lengthy,

"dry" lecture perhaps - would not be sufficiently engaging to hold the attention. In the latter situation little long-term perceptual or conceptual learning would take place, although the student might be able to cram in enough facts to pass a written test in class.

6 Conclusions

"The more abstract the truth one wishes to teach, the more must the senses first be seduced towards it."
Nietzsche (1882)

Education had traditionally erred on the side of absence, by emphasising the conceptual at the expense of the perceptual. VR, on the other hand and by its nature, tends towards the opposite imbalance, emphasising the perceptual at the expense of the conceptual. Perceptual seduction arouses the desire to learn, and students find VR highly motivating. This, along with its flexible but also concrete character, give VR enormous potential in educational settings. In some ways this is also a subversive potential, since it diminishes the distinction between real-life learning and creativity, and what goes on in the classroom. It fosters a style of learning that is less explicit, but more robust, than conventional classroom teaching.

In VR, as in life, learning and creativity are inextricably bound together. To learn through VR, we engender presence (a sense of being there, emphasising concrete perceptual processing), but we must also provoke absence (being immersed in thought, with an emphasis on abstract conceptual processing). These two sides of how people deal with information are needed to respond creatively to what we find around us, whether real or virtual, and whether taught or discovered. Learning is creativity, creativity is learning.

7 Acknowledgements

David Modjeska created the worlds (with a little help from his friends) and carried out the experiments described in Section 3, whilst on secondment from the University of Toronto. We thank him for his many 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.
- Boden, M. A. (1995). Creativity and Unpredictability. In *Constructions of the Mind: Artificial Intelligence and the Humanities*. *Stanford Electronic Humanities Review*, 4 (2).
- Bystrom, K.-E., Barfield, W. and Hendrix, C. (1999). A Conceptual Model of the Sense of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 8, (2) 241-244.
- 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.
- Fencott, P. C. (1999). Presence and the Content of Virtual Environments. Presented at the *2nd International Workshop on Presence*, University of Essex, UK, April 1999.
- 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.
- 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 D. E., and Bovair S.S. (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.
- Lindh, E. (1997). Sensational Interfaces: Function and Passion. *STIMDI Conference Proceedings*. Linköping University, September 1997, pp41-44.
- Lindh, E. and Waterworth, J. A. (1999). To be or not to be conscious. In Ascott, R, (ed.) *Reframing Consciousness*. Exeter, England: Intellect.
- 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
- 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.
- 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.
- Salzman, M. C., Dede, C., Loftin, R. B. and Chen, J. (1999). A Model for Understanding How Virtual Reality Aids Complex Conceptual Learning. *Presence: Teleoperators and Virtual Environments*, 8, (3) 293-316.
- Schubert, T., Friedman, F. and Regenbrecht, H. (1999). Decomposing the Sense of Presence: Factor Analytic insights. Presented at the *2nd International Workshop on Presence*, University of Essex, UK, April 1999.
- Schär, S. G. (1996). The influence of the user interface on solving well- and ill-defined problems. *International Journal of Human-Computer Studies*, 44, 1.18.
- Slater, M., Usoh, M. and Steed, A. (1994). Depth of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 3, (2) 130-144.
- Stappers; P. J., Flach, J. M. and Voorhorst, F. A. (1999). Critical Ratios as Behavioural Indices of Presence. Presented at the *2nd International Workshop on Presence*, University of Essex, UK, April 1999.
- Tulving E., Pearlstone Z (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 5, 381-391
- Tversky, R. A. (1980). *Creativity and Psychopathology*. New York: Praeyer.
- 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 (1996a). VR for Animals. Presented at *Ciber@RT'96*, Valencia, Spain, November 1996.
- Waterworth, J. A. (1996b), 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. (1997a). 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. (1997b). Creativity and Sensation: The Case for Synaesthetic Media. *Leonardo*, 30 (4), 1997, 327-330.
- Waterworth, J. A. (1999a). Technology in Support of Returning - From Conscious Doing to Consciously Being. In Richard L. Amoroso, Rui Antunes, Claudia Coelho,

- Miguel Farias, Ana Leite and Pedro Soares (eds.) *Science and the Primacy of Consciousness*. Orinda, California, USA: The Noetic Press, 1999.
- Waterworth, (1999b). The Scope of Presence. Presented at the *2nd International Workshop on Presence*, University of Essex, UK, April 1999
- Waterworth, J. A. and Chignell, M H (1991). A Model of Information Exploration. *Hypermedia*, 3 (1) 35-58.
- Waterworth, J. A. and Lindh Waterworth, E. (1999). Education as Exploration: being, going and changing the world. *4th Panhellenic Conference on Didactics of Mathematics and Informatics in Education*, University of Crete, Greece, October 1999.
- Whitelock, D. and Jelfs, A. (1999). Understanding the Role of Presence in Virtual Learning Environments. Presented at the *2nd International Workshop on Presence*, University of Essex, UK, April 1999.
- Witmer, B. G., & Singer, M. J. (1994). *Measuring Immersion in Virtual Environments* (Tech. Report 1014). Alexandria, VA, U.S.A: Army Research Institute for the Behavioral and Social Sciences.
- Yates, F. A. (1984). *The Art of Memory*. London: Routledge and Kegan Paul. First published in 1966.
- Zahoric, P. and Jenison, R. L. (1998). Presence as being-in-the-world. *Presence: Teleoperators and Virtual Environments*, 7, (1) 78-89.