Special considerations for navigation and interaction in virtual environments for people with brain injury

A Lindén[1], R C Davies[2], K Boschian[1], U Minör[1], R Olsson[2],
B Sonesson[1], M Wallergård[2] and G Johansson[2]

[1]Department of Rehabilitation, Lund University Hospital, SE 24385 HÖÖR, SWEDEN.

[2]Division of Ergonomics, Department of Design Sciences, University of Lund,
Box 118, SE 22100 LUND, SWEDEN.

*anita.linden@skane.se, roy.c.davies@ieee.org, kerstin.boschian@skane.se,
ulf.minor@skane.se, robert.olsson@design.lth.se, bengt.sonesson@skane.se,
mattias.wallergard@design.lth.se, gerd.johansson@design.lth.se*

ABSTRACT

When a Virtual Environment (VE) is designed, decisions regarding the navigation of the viewpoint, interaction with objects, and the behaviour of the VE itself are made. Each of these can affect the usability and the cognitive load on the user. A VE that had previously been constructed as a prototype tool for the assessment of brain injury has been studied to establish the consequences of such design decisions. Six people, two with brain injury, have used the VE to perform a specific task (brewing coffee) a total of ten times over two sessions separated by a week. These trials were video recorded and analysed. Results and implications are presented and discussed.

1. INTRODUCTION

The use of Virtual Environments (VE) for cognitive rehabilitation is an active area of research and is typified by co-operation between research institutions, hospitals and rehabilitation centres (Riva *et al*, 1999). Application areas are diverse including phobia treatment (for example, Weiderhold and Weiderhold, 1999), treatment of post-traumatic stress disorder (Hodges *et al*, 1999), cognitive assessment (and rehabilitation) (Riva *et al*, 1999), training of people in daily living tasks (Brown *et al*, 1999) and many others. However, the majority of these address the issue of whether a VE can be used for a particular application, and infer from this that the system itself, consisting of hardware and software, is thus adequate.

Virtual Environments supposedly allow people to be within an artificial environment as if in reality. This is achieved through a combination of special hardware for capturing human movements, speedy computer calculations and hardware for displaying the result to human senses. Much philosophical and physiological work has been performed in the area of increasing the level of presence – the feeling of being in a VE, though an in-depth discussion of this is beyond the scope of this paper. However, the experience of a VE as yet is far from being mistaken for reality. One of the biggest hurdles is the devices that a person is required to use for imposing their will on the system and for perceiving the result thereof.

When a VE is designed, decisions are made as to the hardware platform, software and peripherals, as well as the actual content and functionality of the VE itself. Each of these can have a dramatic effect on the usability. In the area of cognitive rehabilitation, it is essential that a VE tool is not rendered unusable due to early design decisions, particularly as people with a cognitive disability may be less tolerant to a poor interface.

To this end, we have investigated the effects of the design decisions on a VE built for a PC desktop VR system, originally as a prototype for cognitive assessment. This paper details the construction of the VE including guiding principles used and assumptions made, describes the studies performed, the observations made, and the implications of the original decisions.

2.   PROJECT DESCRIPTION

The Department of Rehabilitation at Lund University Hospital and the Department of Ergonomics at Lund University are presently co-operating in a long-term project with the following goals:

- to determine whether a VE tool can be a useful complement for the rehabilitation of people with an acquired brain injury and an effective tool in everyday life;
- to find the optimal interface between a VE and the user (and to determine which groups of people with brain injury would be able to use a VE);
- to investigate transfer of training of practical tasks learnt using a VE to the real world; and
- to develop at least three practical applications of VE for rehabilitation.

Each department brings its own unique areas of competence to the project; the Department of Ergonomics is specialised in human computer interaction, and the development of VEs for various applications, whilst the Department of Rehabilitation has expertise in the practical and theoretical aspects of rehabilitation of people with brain injury. The latter was also the first clinic in Europe to be acknowledged by the Committee for Accreditation of Rehabilitation Facilities (CARF).

2.1 The user group

A large group of people who suffer a brain injury are able to retrain their daily living skills and again participate in society. This group is not homogeneous and includes those who have suffered a brain injury through either illness or trauma. Furthermore, cognitive difficulties can be both diverse and obscure. The VE based rehabilitation tools are therefore aimed at this group as it is hoped they can most benefit. In order to reduce complexity, is initially assumed that the people are able to physically manage ordinary computer equipment such as buttons, touch-screen, mouse and joystick, though we do not assume dexterity in both hands.

2.2 Previous work

To date, we have looked at the potential of using a VE for brain injury rehabilitation through testing and interviews with occupational therapists (Davies *et al*, 1999). A simple environment was built based on a well documented, standard, real world assessment of cognitive function used currently by occupational therapists – the task of brewing coffee. The results and comments indicated that the tool has potential as a complement to existing methods of assessment and training, particularly at an early stage when the patient is physically unable to perform the task in reality, in situations where it isn't practical to train repetitively in reality, and for people who may be more interested in computers than brewing coffee. Many suggestions were also made as to other applications of VE technology to brain injury rehabilitation. It was concluded beneficial, therefore, to continue and to delve deeper into what happens at the interface between the user and the computer when sharing a VE.

3.   HUMAN – VE INTERACTION

There are many aspects to the human-VE interaction problem, such as:

- physical loading problems associated with wearing often heavy devices or holding limbs in uncomfortable positions for a long period (Nichols, 1999);
- physiological effects such as nausea and eye discomfort when using a HMD (Cobb *et al*, 1999);
- method of interaction with the VE; and
- extraneous cognitive load – that is the load on the user above and beyond that of performing the task which can be attributed to the usage of the tool itself.

The first of these can be mostly side-stepped by using desktop VR with standard input devices, though this may increase the cognitive load due to a lesser degree of immersive feeling. The last two are the basis for our further study.

Cognitive load is an all-pervasive factor, which must be considered at all stages of VE design and usage. A general rule-of-thumb is that too much extraneous cognitive load will distract from the concentration on the task (as[one must instead concentrate on just interacting with the VE). People who have cognitive disabilities may be especially sensitive to this, so it is important to try to reduce this effect as much as possible.

The method of interaction can be further broken into navigation of oneself in the VE, and interaction with objects in the VE.

3.1 Navigation of oneself in the VE

Navigation in a VE can have two meanings; finding one's way in a large environment; or manipulating the viewpoint (as seen through the computer screen) to see what one wishes to. The latter is most directly connected to the input devices and is of interest to this study, though the former does also have an effect, particularly for people who find it difficult to remember parts of the VE not currently in view.

The form of navigation can range from completely automatic to self-controlled, with various forms of 'half-automatic' in between. Navigation can be either from a first person view – you see on the screen what you would see from your own eyes in the VE; or third person view – you control an avatar which you see perhaps from behind.

Automatic navigation implies that the computer tries to decide what the user wishes to look at depending on their intention. Their intention is inferred from their input and previous actions (and thus the state of the system). The user input is used to initiate events, hence navigation of the viewpoint comes as a side-effect of the initiation of events (of interaction). This form of navigation is expected to be least cognitively taxing on the users, since the only interaction with the computer is to initiate events, which is directly related to the performance of the task. If the user also has to position the viewpoint, this is an extra level of complexity that one normally doesn't have to worry about in reality (do you notice that you move from place to place when making coffee? (Perhaps if movement is difficult, for example in a wheelchair, but otherwise, probably not.)

Self-controlled navigation is performed through some input device such as a joystick, mouse or keyboard. The type of device and how it is programmed can affect the usability of the device, for example, Peterson *et al* (1998) have found that a joystick allows precise manoeuvring, but that a Virtual Motion Controller (a type of stand-on-platform joystick) is better for route learning. It is usual to limit the number of degrees of freedom to that required by the application. There are two basic flavours of self-controlled navigation; walk through and fly through. In the former case, the height is fixed and two degrees of freedom are permitted; movement forwards, backwards and turning to the left or right. Sometimes an extra degree of freedom is added to allow sideways movement. One can also normally walk up and down stairs, and perhaps jump onto low objects. Turning can be as for a person, in that the view can be spun on a point without forward or back movement, or it can be as for a car in that one must also move forward or back to actually turn, and the turn is an arc.

Automatic navigation is best suited to situations where the entire VE can be viewed at once, whereas, some form of self-controlled navigation is required when the VE is large or has hidden areas.

In between these extremes, there is what may be dubbed *half-automatic navigation*. This is where the user has a level of control, but allows the computer to aid in an intelligent way. Half-automatic can use a variety of algorithms, for example, with the user still controlling navigation with an input device, the computer always moves the viewpoint towards and away from the last selected object (or perhaps the object over which the mouse pointer is currently resting).

Finally, there is a further complication of body versus head movement (coarse versus fine movement). Many 3D games use a combination of mouse movement for head positioning, for example, and keys or joystick for body positioning. However, this requires the user to be adept with several input devices at once.

3.2 Interaction with objects within the VE

Once the user has positioned the viewpoint, they may wish to actually do something (since there is a task to be performed). There are at least three things one can do with virtual objects (or indeed real objects):

1. *activate objects* in various ways such as turning on a switch, opening a packet of coffee, pushing a button or turning off a tap;

2. *move objects* from one place to another (and rotate if appropriate); and

3. use one object with another (*object-object interplay*), for example, using a knife to spread butter on bread, or using a coffee scoop to take coffee from a packet to put into a filter.

The user tells the computer to do one of these by giving a command through an input device. This in turn may initiate an event or cause a change in the state of the VE. There is a plethora of input devices such as: standard PC mouse, dataglove, gesture recognition, voice recognition, force feedback haptic device, touch screen and interaction with a real object to effect a change in the VE. Some of these double as display devices. Again, the input device itself affects the usability. Werkhoven and Groen (1998), for example, performed a study comparing object manipulation performance using a dataglove and a six degree of freedom spacemouse in an immersive VE. This showed that a dataglove wins in both accuracy and speed for such tasks. However, similar work has not been found with regard to desktop VE input devices.

4. A STUDY OF HUMAN – VE INTERACTION.

In the design of any VE, a number of sometimes arbitrary decisions must be made concerning the computer system, the input and output devices, the structure of the interaction and of the VE itself. These affect the usability of the VE. The aim of this study was to take a deeper look at the effects of the design decisions on the coffee making VE assessment tool. These are grouped into overall guiding principles which were used whenever a decision was to be made regarding the design, and general assumptions.

4.1 Guiding Principles

The guiding principles were primarily based on a wish to reduce the complexity of the interaction with the VE and thus reduce the extraneous cognitive load on the user. At a later stage, these may also be questioned.

1. Only things that can exist in the real environment should be in the virtual, for example, no extra buttons or icons such as found in ordinary interface designs. Similarly, objects should act as expected in the real environment. *Reason:* People are accustomed to the real world and tend to notice when objects in a VE differ from reality. Furthermore, keeping the VE 'pure' avoids the question of what extra features are required and the effect on the user, both in cognitive load and effect on transfer of training.

2. The user must have free choice in the initiation of events, being able to choose to do things in whatever order they wish – even the wrong order (this implies that the system must react in a sensible way in all situations, with regard to principle one). *Reason*: One way of learning is by trial and error. Similarly when the VE tool is to be used for assessment, it is the errors that give the most information. Therefore, all (or at least almost all) possible orders of task performance must be possible and error situations programmed in (e.g. water running out of the coffee machine if the pot is not returned).

3. Where a decision has to be made between two or more design alternatives, the simplest should be chosen. The definition of 'simplest' may vary from case to case but must always take into regard cognitive load for the user. *Reason*: We are striving for a low extraneous cognitive load. However, in many situations, it may not always be clear which choice will be better, therefore, careful documentation and evaluation is essential.

4.2 Assumptions

In the making of the coffee machine VE, the following assumptions were made:

4.2.1 Hardware and Software. A standard PC was used, not a high-end graphics machine. This was due mainly to cost and availability of such machinery in the hospital environment. Furthermore, modern graphics cards for 3D games are now capable of rendering at a more than adequate rate, and stereo sound cards are both common and cheap. The software used was Superscape VRT, and the VE included realistic sound effects.

Similarly, standard input and output devices were assumed, in this case, mouse, computer monitor and speakers (desktop VR). This was again due to cost, compatibility with ordinary computers and availability, but also since the effect of putting a head mounted display on somebody who has suffered a brain injury is as yet uncertain. Furthermore, people can learn well using desktop VR. (Brown *et al*, 1999). No device was required for movement of the viewpoint as this occurred automatically.
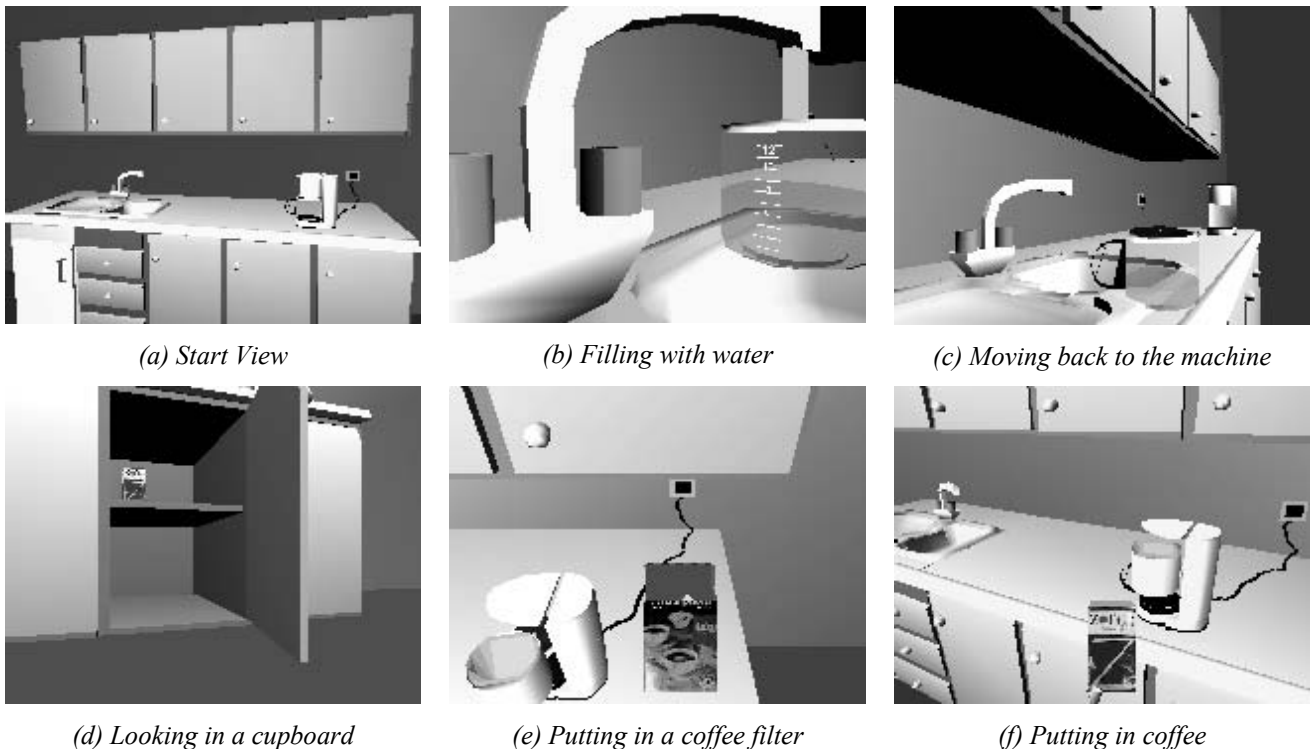
4.2.2 Navigation of the viewpoint. Navigation was completely automatic.
First person view was used to bring the user conceptually closer to the VE – it is themselves performing a task – thus perhaps avoiding an extra level of complexity. As a bonus, we also avoided the problem of a virtual body covering parts of the already cramped viewing area.

4.2.3 Interaction with virtual objects. All object interaction was by single mouse clicks (no drag-and-drop, double clicking or other complex movements were required).

▪ Objects could be activated by single mouse clicks (requiring no turning, dragging or pulling actions).

▪ Objects could be moved by 'picking up' which placed them in the foreground (Fig. 1 c, e, f). This was performed by clicking on the object. Objects thus picked up were then carried (without a visible hand) with the viewpoint and could subsequently be put down by clicking where they were to be placed.

▪ An object could be acted upon by another object by using the object in the hand with the next object clicked upon (Fig. 1 e, f)

Objects which could be activated, picked up, put down or acted upon by another object didn't display any feedback before being clicked upon (such as mouse pointer alteration, sound cues or highlighting). This was due to principle 1 above.

(a) Start View *(b) Filling with water* *(c) Moving back to the machine*

(d) Looking in a cupboard *(e) Putting in a coffee filter* *(f) Putting in coffee*

Figure 1. *Various stages in the coffee making process.*

5. METHOD

To establish the effect of the design decisions and assumptions on the usability of the coffee brewing VE, a series of case studies were performed in which six people were instructed to brew coffee using the VE tool. The subjects repeated the task a total of ten times over two sessions separated by a week to allow for an estimation of short and long term learning effects. Each subject used the VE singly and had no previous experience of it.

5.1 The Subjects

The subjects consisted of two patients with acquired brain injury (one man and one woman) and four hospital personnel (one man and three women), all with word-processing level computer experience, though not 3D game players. A couple had used a computer a little more in their daily work. The ages ranged from 35 to 58. All could make coffee in reality.

5.2 Experimental Set-up

To record material for analysis, two video cameras were used, one at the front of the subject to capture facial expressions, body movements and sound, and one from the side to capture the whole scene as well as input device usage. A video signal was also taken from the computer and mixed in with one of the other video signals.

At the start of the first session, each subject was shown a video that illustrated how to use the VE. They were then asked to 'brew coffee' using the VE, which they repeated five times. A week later, each subject came back and was again asked to brew coffee five times using the VE. This time, there was no instruction beforehand. At the end of the second session, an interview was conducted. This consisted of 14 questions to which the subjects could freely answer. Questions were aimed at establishing the subjects' own views of the tool, thoughts on the navigation and interaction methods and to highlight any particular problems experienced. A couple of questions were also aimed at determining the experience of the test process itself.

5.3 Analysis Methods

The analysis was performed from the videos, observations, and the interview. The coffee brewing task was broken logically into ten steps which are categorised according to object interaction method (Table 1).

From the videos, a detailed description of the actions, specific problems, comments, and other signs from the subject were made for each step. Times were gauged for the performance of each step, for each trial in

each session for each subject. Total completion times were measured from the videos and an optimal time was computed for each step as the median value of each step from the six fastest total times. The optimal times were then used to normalise the data and to calculate the normalised median value for each step over all the subjects and tests. A tally was also made of the number of instances the time to perform a step was more than double the optimal time. Median values were used rather than averages to reduce the effect of large outliers. Comments from the interview were summarised.

Table 1. *Coffee brewing steps.*

Step	Task Description	Interaction Method
1	Take the coffee pot from the machine to the sink.	object movement
2	Fill the coffee pot with water.	object activation
3	Take the coffee pot to the machine, put the water in and close the lid.	object movement + object activation
4	Put the coffee pot on the hotplate in the machine.	object movement
5	Take the coffee filters from the cupboard and place one in the filter holder.	complex
6	Put away the filters.	object movement + object activation
7	Take the coffee from the cupboard and put scoops of coffee into the filter in the holder.	complex
8	Put away the coffee.	object movement + object activation
9	Close the filter holder.	object activation
10	Turn on the machine.	object activation

Note that viewpoint movement is not explicitly mentioned, as this is automatic. Steps 5 and 7 include several object interactions, though primarily object-object interplay.

6. RESULTS

6.1 Videos and interviews

Specific points of note from the analysis of the videos and the interview are summarised below.

Automatic Navigation, for the most part, seemed to present no major problems. However, some situations occurred where the view made it difficult to see clickable objects. A need for a means to change the view was apparent, mainly to take a step back. Five of the subjects tried to find something to click on to change the view and one expressed a need to get closer to objects. One subject mentioned that the view changing was considered a confirmation of being on the right track, though one of the patients found the automatic view changing to be confusing.

Object Activation was also managed without major problem. The main difficulty was in finding the sensitive area for clicking on, with some people clicking just beside objects to active, particularly if the current view made the object small. One subject expressed a need for quick feedback of correct object activation.

Moving Objects provided some problems, particularly in the start as the subjects tended to try to drag-and-drop the objects. This was even the case for one subject who expressed no knowledge of drag-and-drop. The sensitive areas tended to be missed again, and all of the subjects had trouble comprehending when objects were being held (though two assured that the concept was understood). The sensitive areas for the placement of the coffee pot on the coffee machine element or the sink were often missed.

Object-Object interplay posed one specific problem. When the coffee filter holder was open with a filter placed inside, and the packet of coffee was being carried, clicking on the filter holder took over a scope of coffee grounds to the filter. All but one of the subjects had difficulties figuring out how to tell the system when there was enough coffee and to close the filter holder.

The subjects showed a tendency to be disturbed by their own mental models, both in the expectation of

how the coffee machine worked and in how the VE worked. One was used to another type of coffee machine and found it difficult to get used to this one. This person also insisted in replacing the coffee packet in an upper cupboard though every time it must be taken from a lower cupboard. Another subject insisted that all the cupboard doors must be closed to continue. Yet another subject wanted to try to remove the filter from the holder (it actually comes out automatically), lift the lid of the filter holder (rather than rotate out the filter holder) and turn on the water to initiate the action of moving the coffee pot to the sink. All subjects showed a reluctance to let go of these mental models even in the face of contradictory information.

In terms of the experience of using the VE, it was considered to require some concentration, particularly in the beginning; some subjects tended to forget to perform certain actions such as closing the filter holder and putting the water in the machine but remembered after a few trials of finding water on the floor or no coffee in the pot; and only one of the subjects (a patient) considered being video filmed as disturbing

Finally, one of the patients showed a tendency to always click slightly to the right of objects – this might have been due to a slight degree of neglect or a visual problem. The other patient required objects to always be on the right of the coffee machine (which was not possible) so spent quite a bit of time trying to place them there.
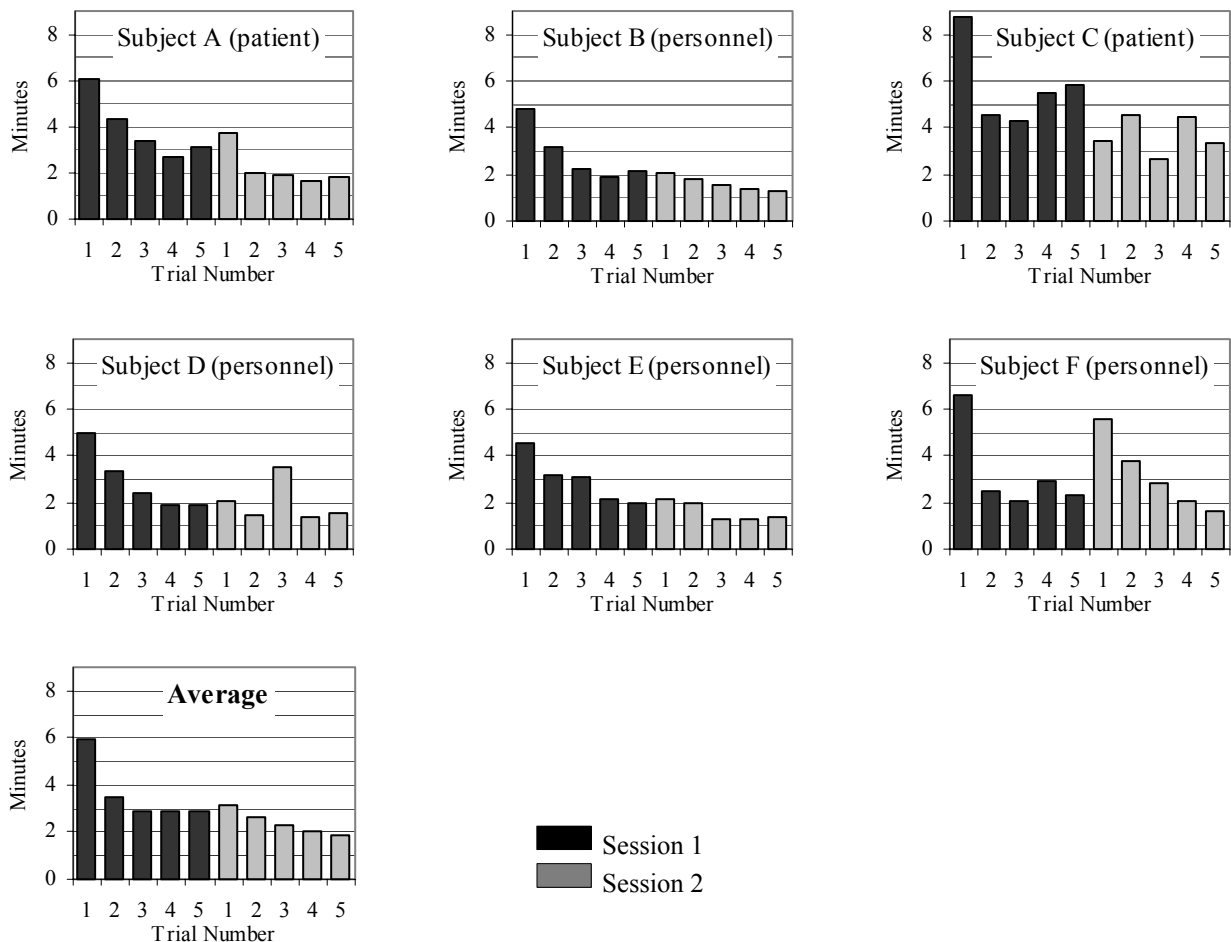
6.2 Total Times

Figure 2. *Total times taken to complete the task for each subject.*

6.3 Comparison with optimal times

Comparing with optimal times (Table 2), the normalised median values for the object activation steps are close to a value of 1, meaning the there was little difference from the optimal times. The percentages of datapoints greater than twice the optimal are also low, signifying that most were less than twice as slow.

For the move object steps, the medians are higher, signifying an increase in time compared to the optimal, and the percentages of datapoints greater than twice the optimal are in some cases over 50%. For the more complex steps, including object-object interaction, the medians are also high, averaging to a time to perform the steps of just under twice that of the optimal times. Almost half the datapoints are greater than twice the optimal.

Table 2: *Normalised medians for task performance of each step*

Interaction Method	Activate			Move		Move + Activate			Complex	
Step Number	2	9	10	1	4	3	6	8	5	7
Normalised Median	0.78	1.00	1.00	2.00	2.00	1.38	1.57	1.33	1.52	2.13
Percentage of datapoints greater than twice optimal	7	20	7	52	51	18	27	25	30	58

7. DISCUSSION

Automatic Navigation appeared to provide no undue problems to the subjects, apart from when a viewpoint position was encountered which made it difficult to click on the next desired object. The main difficulty was when a position was taken close to the bench, but the subject wanted to back away and take a whole view again. A 'back' button might help for this. However, automatic navigation is not always possible whilst still allowing the user free choice in event initiation. For example, If the VE is sufficiently complex so that there are parts that are not visible at all times, then principle 2 is contradicted in that the user doesn't have free choice to initiate events in areas that are not visible, if they cannot navigate themselves there.

Few problems were noted with object activation. Subjects seemed not to be worried that a 'click' was used in situations where in reality one should grasp and turn (such as a tap), or grasp and lift, or push down, as for a switch. In all these situations, the object to be activated appeared the same as in reality and provided immediate feedback of success (such as the tap turning and the water running). The normalised medians showed, in fact, an improvement over the optimal values (since the six best completion times did not necessarily mean that all steps were completed quickly) and most of the subjects completed object activations in a time well less than twice the optimal.

Moving objects, on the other hand, seemed considerably more complicated. The normalised medians show that the time to move objects was almost double that of the optimal times and that steps one and four caused most problems with over half the values being over twice the optimal times. In the videos, it could be seen that the subjects had trouble knowing where to click for step one to move the coffee pot to the sink and step four when the coffee pot should be set on the hotplate, both problems with the sensitive areas. In step four, the subjects often clicked in the area of the coffee machine, however, its shape meant that the centre is actually a gap, and a click there goes through to the bench. The result was that the coffee pot was put down on the bench instead of on the hotplate. To avoid such problems, invisible sensitive areas could be put around such items so that a click in the general vicinity is sufficient. However, sensitive areas cannot overlap, as this would cause confusion.

Oddly enough, all the subjects wanted to drag-and-drop objects around the VE in the beginning, even though this wasn't possible. Whether this comes from having some computer skill, or whether this behaviour is natural cannot be determined, but would make an interesting further study. One subject maintained they had some computer skill, but not of drag-and-drop, though it is difficult to see how that could be the case. If it can be said that drag-and-drop is a natural behaviour and that all people, even those with no computer skill, can manage it, maybe it could be used in the construction of VEs for people with brain injury. All the subjects misunderstood the concept of carrying objects, possibly since it wasn't quite clear that the object in front of the view was currently being 'held'. Perhaps if a virtual hand had grasped the object, this confusion wouldn't have occurred.

Steps three, six and eight, while also requiring objects to be moved were less problematic than steps one and four. Maybe the sensitive areas in the cupboards were more easily fathomed.

Object to object interplay caused some confusion, the main problem resulting from the single-click assumption. With this, it is possible for situations to occur when a click can be interpreted in several ways. This was most apparent when the subjects were putting coffee into the filter. There was little problem in clicking on the filter holder to mean 'put a scoop of coffee here', however, when sufficient coffee had been placed, it was then logical to click on the filter holder to close it. The problem, however, was how to let the computer know this change in mind. Adding the capability for drag-and-drop might help. In this case, coffee could be put into the filter by dragging the scoop to the filter, and the holder closed by a single click. Using the right mouse button for extra functions, though is not desirable as this would complicate the input device considerably.

Considering the total time graphs of Fig. 2, all the subjects had the longest time for the first trial of session one, not surprisingly since they had never used the VE before, then improved for each subsequent

trial. In some cases, times started to increase again near the end of session one, perhaps due to fatigue (or boredom). From the analyses of the videos, trial one included many instances of the subjects attempting to understand how to make objects do what they want. When they finally succeeded, this lesson was carried over to the next trial. Subject C, however, was unable to accept that objects could not be placed to the right of the coffee machine and attempted to do this in every trial, despite previous attempts and eventually having to give up.

In session two, all the subjects except subject C showed an increase in time compared to the last trial of session one, suggesting a certain amount of forgetting of what was learnt previously. On average, there seems to be both short term and long term learning effects occurring. The implication is that even if the interface is not perfect, people will learn to compensate, assuming the capacity for learning is not impaired. Subject C, however, demonstrates that even simple things like not being able to place objects to the right of a coffee machine may disrupt the concentration on the task.

Interestingly, nobody complained about reality glitches due to programming VE limitations such as unrealistic water behaviour. It seemed sufficient that the end states for each action were correct, rather than how objects interacted to get there. So, for example, the water from the tap doesn't meet the surface of the water in the pot when filling, but water appears there anyway, and the pouring of the water into the machine occurs without seeing the water go between. In these cases, it might be however, the sound effects that fill the visual gap. Subjects had more problem with remembering that the coffee machine was of a different model to that they were used to.

8. CONCLUSION

In the design of a VE system for use by people with a brain injury, it is essential to carefully consider all aspects of the system design, from the input and output devices to the contents and structure of the VE itself. Every aspect can potentially have a negative effect on the cognitive abilities of the user, which for a person with brain injury may make the VE unusable.

With this in mind, we have evaluated the effect of design decisions for a VE system intended as a complement to existing techniques for assessment of cognitive function. The following conclusions can be made:

- People seem to have an inherent understanding for click-to-activate and drag-and-drop (this was the case even for the two subject with brain injuries).
- Moving objects poses some problems though, mainly in choosing where to click the mouse to place the object. The users' mental models of real world objects appears to aid in deciding which objects afford being clicked upon in the VE for activation. However, for object placement, the mouse-click means "place there", an apparently less logical concept. Perhaps some visual cue may help.
- The automatic navigation technique works well for situations where the VE is not too large to be viewed in one screen-full, though some situations can occur when a key would be useful for returning to some overview position. As navigation was hardly noticed by the subjects, it would be ideal for VEs to be used by people with a brain injury, though not in too large environments.
- Problems can occur when a mouse-click can be interpreted by the computer in several ways, though it may be possible to resolve these with clever programming and alternative interaction metaphors.
- A certain amount of imprecision must be allowed in the sensitive areas for object interaction. For example, when an object is small on the screen, being able to click in its general vicinity should assist the interaction process. Similarly, people seem to like to click in the middle of objects, so if there is a hole there, maybe an invisible sensitive area would help.
- When carrying objects, it is essential that it is clear that the object is being carried, perhaps by having a hand holding it.
- The users' mental models can affect the usage of the VE. Firstly, disparities between virtual objects and real ones the user is accustomed to can cause misunderstandings. Secondly, expectations in how the VE itself works, if wrongly made, can cause confusion. The main problem is when actions allowed in reality are not permissible in the VE.
- Persistent mental and physical problems of the users were apparent from the analysis of the videos. These showed, in one case, a rigidity of mental model and a possible physical problem in another. Therefore, using such a VE in this manner, recording the users' actions and screen view, and performing a detailed analysis could be developed into a workable assessment tool for brain injury rehabilitation.

- The patients and one of the personnel showed a tendency to become tired after performing the same task a few times. However, all the subjects managed to perform the task within acceptable time limits, thus showing that a VE tool could well be usable for rehabilitation.

Further work is planned to look into other forms of viewpoint navigation and the effect of the input device on usability as well as transfer of training effects and the development of a number of complementary VE applications for brain injury rehabilitation. For navigation, the interplay between the number of degrees of freedom and the type of input device will be investigated. For object interaction, the use of drag-and-drop and whether a virtual hand to hold objects will simplify object movement will be investigated.

9. REFERENCES

Brown, D., Neale, H. and Cobb, S. (1999) Development and Evaluation of the Virtual City, *International Journal of Virtual Reality,* **4**, *1*, pp 28-40.

Cobb, S. V. G., Nichols, S., Ramsey, A. and Wilson, J. R., (1999), Virtual Reality-Induced Symptoms and Effects (VRISE), Presence, **8**, *2*, pp 169-186.

Davies, R. C., Johansson, G., Boschian, K., Lindén, A., Minor, U., Sonesson, B. (1999), A practical example using Virtual Reality in the assessment of brain injury, *International Journal of Virtual Reality,* **4**, *1*, pp 3-10.

Hodges, L. F., Olasoc Rothbaum, B., Alarcon, R., Ready, D., Shahar, F., Graap, K., Pair, J., Hebert, P., Gotz, D., Wills, B., and Baltzell, D., (1999), A Virtual Environment for the Treatment of Chronic Combat-Related Post-Traumatic Stress Disorder, Clinical Observations During Virtual Reality Therapy for Specific Phobias, *CyberPsychology and Behavior*, 1999, **2**, *2*.

Nichols, 1999, Physical Ergonomics Issues for Virtual Environment Use, Applied Ergonomics, 30, 79-90, 1999.

Peterson, B., Wells, M., Furness III, T. A. and Hunt, E., (1998), The effects of the interface on navigation in virtual environments, *Proceedings of the Human Factors and Ergonomics Society 1998 Annual Meeting*, pp 1496-1505.

Riva, G., Rizzo, A., Ph.D., Alpini, D., Barbieri, E., Bertella, B., Davies, R. C., Gamberini, L., Johansson, G., Katz, N., Marchi, S., Mendozzi, L., Molinari, E., Pugnetti, L., Weiss, P. L., (1999), Virtual Environments in the Diagnosis, Prevention, and Intervention of Age-Related Diseases: A Review of VR Scenarios Proposed in the EC VETERAN Project, *CyberPsychology and Behavior*, 1999, **2**, *6* pp. 577-592

Weiderhold, B. K. and Wiederhold, M. D., (1999), Clinical Observations During Virtual Reality Therapy for Specific Phobias, *CyberPsychology and Behavior*, 1999, **2**, *2*.

Werkhoven, P. J. and Groen, J., (1998), Manipulation Performance in Interactive Virtual Environments, *Human Factors*, **40**, *3*, pp 432-442