Exploration and Virtual Camera Control in Virtual Three Dimensional Environments

Colin Ware and Steven Osborne
Faculty of Computer Science
University of New Brunswick
P.O. Box 4400, Fredericton
New Brunswick, E3B 5A3,Canada.

Abstract

This paper evaluates three distinct metaphors for exploration and virtual camera control in virtual environments using a six degree of freedom input device. The metaphors are "eyeball in hand", "scene in hand", and "flying vehicle control". These metaphors have been implemented and evaluated using an IRIS workstation and a Polhemus 3Space. The system has the capability to record the motion path followed during an exploration session and this can be recorded and played back to create a movie. Evaluation is through intensive structured interview sessions wherein subjects are required to complete a number of tasks involving three different "toy" environments. None of the metaphors is judged the best in all situations, rather the different metaphors each have advantages and disadvantages depending on the particular task. For example, "scene in hand" is judged to be good for manipulating closed objects, but is not good for moving through an interior; whereas "flying vehicle control" is judged the best for navigating through the interior, but is poor for moving around a closed object.

Introduction

The problem of providing a user interface for exploring virtual graphical environments is closely related to the problem of viewpoint specification and also to the problem of providing a camera path through a virtual environment. Before describing our approach to the problems of navigation and virtual camera control we briefly review some of the other methods which exist with particular attention to the user interface metaphors they employ.

First we make some observations on the nature of the task.

- The task of placing a viewpoint in a virtual 3D environment has inherently six degrees of freedom — three for positional placement and three for angular placement. An additonal degree of freedom has to be introduced to provide a field of view scale factor, the function equivalent to the zoom on a camera. In this paper we ignore the problem presented by this seventh variable.
- The task of exploring a virtual environment can be accomplished by navigating a viewpoint through the environment. If the motion path is recorded for later playback, the exploration path becomes a virtual camera path.
- Moving a viewpoint around in a virtual environment is isomorphic with moving the environment around in terms of the resulting image. Of course, the user interface metaphor is quite different.

A number of studies have investigated the problem of object placement by extending 2D devices to cover 6D manipulations. Usually translations pose no great difficulty, thus getting rotations to work well has received the most attention. Evans et al [6] transformed the input from a digitizing tablet so that x and y motion yielded rotation about y and x respectively, and a circular stirring motion yielded rotation about z. The interaction metaphor here is that of a virtual turntable on which an object is placed. Chen et al [5] refined this metaphor to that of placing the object in a conceptual sphere. Any motion of the tracker is interpreted by the program as if it were a physical gesture on the surface of the sphere. This allows all three axes of rotation in a single coherent model. The major limitation inherent in using 2D devices is that at least one change of state is required to cover all translations and rotations.

At the other end of the technology spectrum is the creation of so called "cyberspace" which places the user directly in the graphical environment. This work, pioneered by Sutherland [12] involves the use of head mounted monitors which, when combined with head mounted sensors, give the user the impression that he or she can freely turn their head to gaze at features in the graphics scene. A more recent version of this approach [7] includes the "data glove" which makes it possible to encode the position of each of the subject's fingers and hence to "pick up" graphical objects [16]. The user interface metaphor here is that of direct manual manipulation, and it is followed so faithfully so as to seem more literal than metaphorical. However, at present we are perhaps justified in considering this to be a metaphor because the objects have no feel and no mass and the illumination model (where there is one) is primitive. Technical progress may overcome these problems.

Our study of the manipulation problem began with studies using a device which we call the "bat" (because it is like a mouse that flies or *fledermaus*). This device is based on the Polhemus 3Space Isotrac[1], a six degree of freedom spatial sensor which encodes both position and orientation. We have mounted the 3Space sensor in a case shaped like a bar of soap and supplemented it with a single button. In our initial studies we developed an interface which allows this device to be used in "mouse mode", that is, relative changes in position and orientation are measured while the button is depressed which enables a selected object to be ratcheted into position using a sequence of hand movements. We have studied the problems of manipulating a heirarchical three-dimensional scene both from a conceptual viewpoint [15] and to obtain some empirical data on the speed of performing various tasks [14].

The spatial metaphor used in our original study is of a direct mechanical linkage between the bat and the currently selected object. We call it the "scene in hand" metaphor because of the apparently direct manipulation of the scene which occurs. If the user translates the bat, the corresponding translation occurs for the object. If the user rotates the bat about the its center, the corresponding rotation occurs for the object about the object's own center. This is a direct extension of the way in which a mouse is used to move a cursor about the screen; there is a one to one correspondence between mouse and cursor motion. It is also possible to use the "scene in hand" metaphor to change viewpoints; by selecting the root object in the heirarchical scene and twisting the bat. the virtual environment is rotated about the center of the root object.

The above work was successful in showing that it is possible to use the same metaphor for both manipulations and for some sort of navigation. But as we turned our attention more to the problems of exploration and virtual camera control we became aware that at least three very distinct metaphors existed for navigating through virtual environments using the 3Space and a high performance graphics workstation. Initially, we had little intuition about which was likely to be the best. However, before we describe these three interaction metaphors we first address the issue of what the term "metaphor" means in this context.

The Problem of Metaphor

In all of the above discussion we been using the term *metaphor* in the sense of an explicit *simile*. The user understands the behaviour of the interface to be "like" having the world on a turntable or "like" locomoting through the physical environment. These metaphors are provided to give the user an *internal model* [9]. of the interface which will help the user in initially understanding the system's behaviour. This is obviously true for the virtual sphere of [5]. It is equally true, although less obviously so, for the cyberspace environment. The cyberspace metaphor is the most compelling, but it also has limitations (besides those of exorbitant cost). There are the obvious technological limitations relating to realism and lack of touch which may be overcome in time. But what is more interesting is that the metaphor itself creates both affordances (things it is easy to do) and restrictions (things it is hard to do).

User interface metaphors provide two fundamentally different kinds of constraints on their utility. The first of these constraints is essentially cognitive. The metaphor provides the user with a model that enables the prediction of system behaviour given different kinds of input actions. A good metaphor is one that is apt — which matches the system performance well — and also is well understood by the user and easy to learn. One could model an F14 airplane to provide a metaphor for flying the viewpoint around, but if the user does not know how to fly an F14 this may not be very helpful.

The second aspect of a metaphor relates to the physical constraints which its implementation places on the interface. A particular metaphor will naturally make some actions easy to achieve and others difficult. Thus, for example, the cyberspace metaphor places natural restrictions on the acceleration and velocity of the viewpoint and it makes it impossible to place the viewpoint high above the scene because it is limited to positioning the viewpoint where a person can place their head. On the other hand the cyberspace metaphor makes the control of the position of the viewpoint completely straightforward within these limitations.

Here we evaluate three metaphors for moving through virtual environments both in their cognitive aspects and with respect to the sets of viewpoint ma-

[1]3Space is a trademark of Polhemus Inc.

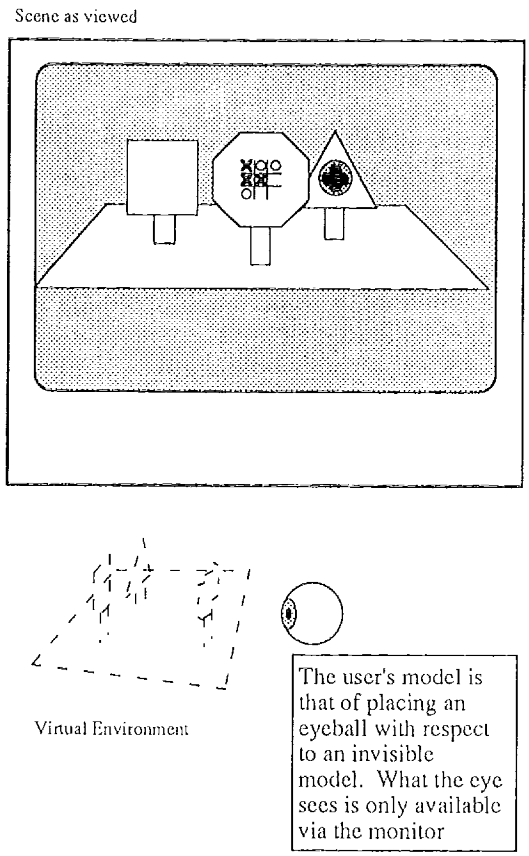

Figure 1. Eyeball in hand metaphor

nipulations that are easy and difficult to achieve.

Interaction Metaphors

In constructing the programs which instantiated each of the interaction metaphors we attempted to keep each metaphor as simple as possible, although interesting elaborations constantly occurred to us. The reason was our desire to study the problems of each in its purest, most essential form. The three interaction metaphors which we implemented are as follows:

Eyeball in hand. This technique involves using the Polhemus as a virtual video camera or eyeball which can be moved about the virtual scene. The virtual scene is like an invisible model which inhabits the room with the user and which can only be seen from the vantage of the hand held eye (see Figure 1). The view from this eyeball is mapped to the screen of the IRIS. This metaphor has previously been studied by Badler et al [1] who found that "consciously calculated activity" was involved, and by Brooks [3, 4] who found it to be useful although he also found that an additional plan view of the scene was useful to prevent disorientation.

Despite our intentions to keep the implementations simple, our first subject encouraged us to add a feature whereby the viewpoint is moved only while the button is depressed. This allows a form of ratcheting: If the user lets go of the button, the view as displayed on the screen, remains fixed. The user can then move the 3Space to a more comfortable position and resume the viewpoint manipulation by depressing the button. Unfortunately, this feature destroys the mental model of a fixed invisible scene. Instead the user must imagine that it is possible to ratchet the invisible scene around using the viewpoint as a kind of handle. This is a confusing mental model, our subsequent experimental subjects preferred not to use this feature, and we do not discuss if further.

Scene in hand. In this metaphor the scene is made to move in correspondence with the 3Space (bat). If the bat is twisted clockwise the scene rotates clockwise; if it is translated left the scene translates left, etc. It is akin to having an invisible mechanical linkage which converts all hand translations and rotations into translations and rotations of the scene (see Figure 2). Large movements of the scene are accomplished by ratcheting which is achieved by means of a button which acts as a clutch. To make a large object

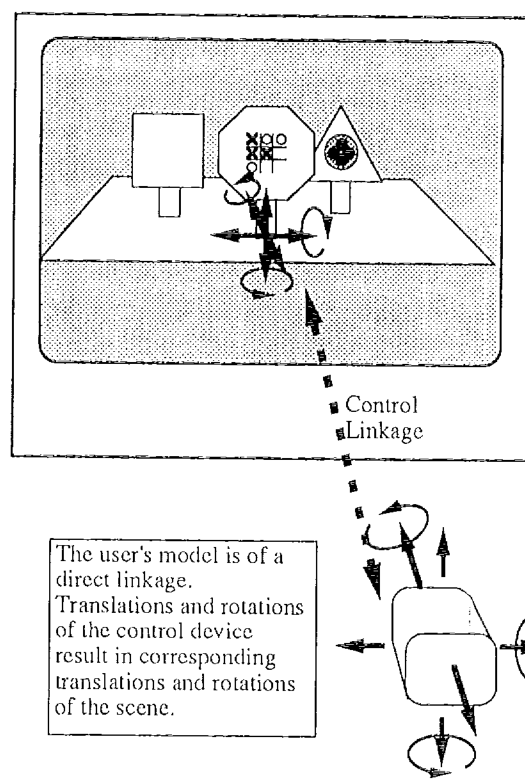

Figure 2. Scene in hand metaphor

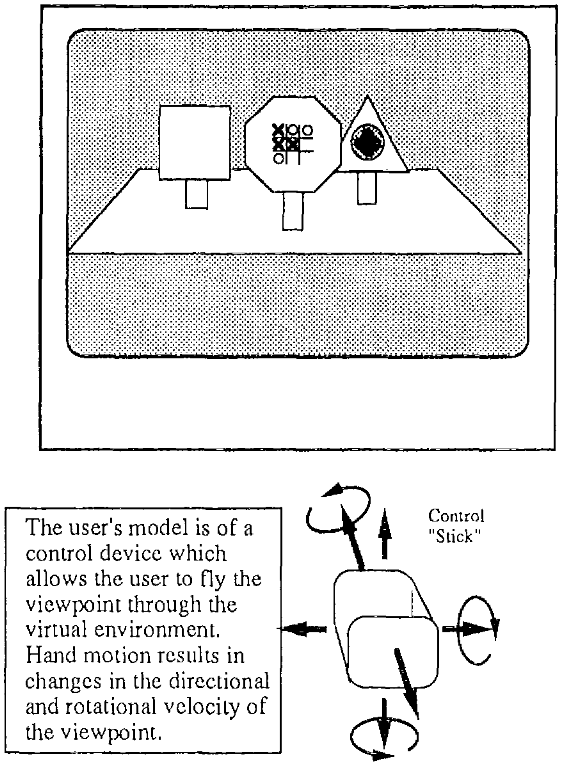

Figure 3. Flying vehicle control metaphor

movement the user depresses the button makes the movement, released the button and returns his/her hand to its start position and repeats this process as many times as necessary.

A previous study in our laboratory has found the scene in hand metaphor to be useful for manipulating discrete objects and changing the viewpoint in a heirarchical scene [15].

A related study by Schmandt [11] implemented an interface in which the user placed a hand in the graphics environment, by using beam splitting mirrors and piezo-electric shutters for stereopsis. The main problem was found to be coping with inconsistent occlusion. The user's hand and the graphical objects both appeared to be semi-transparent — there was no occlusion based on depth, and occlusion is probably the strongest of all depth cues. It was because of these problems that Ware and Jessome decided to implement a one-to-one correspondence, instead of actually placing the user's hand in the scene.

Flying vehicle control. In this metaphor the bat is used as a control device for a virtual vehicle. The virtual environment is perceived from this vehicle. We only call this interface a "flying vehicle" in order to provide the user with a preliminary model of its behaviour. No attempt is made to actually model the characteristics of flight, which would, for example give the user control over acceleration as a method for controlling forward velocity. Instead we control spatial velocity and angular velocity directly (human factors studies suggest that people can in general control velocity far more easily than acceleration [8]. We found it useful to put a non-linear control on translational velocities making the velocity related to the cube of the displacement. This makes it possible to have very fine control while still allowing rapid motion of the viewpoint. Our implementation of the flying metaphor is illustrated in Figure 3. This metaphor has also been studied by [3, 4] who used two joysticks to control velocities. However he fails to mention how many degrees of freedom were under control, neither does he comment on the ease of learning of the metaphor.

Toy Environments

The three motion control metaphors were evaluated in the context of three "toy" environments designed to embody the important properties of different 3D task domains. These environments, which are illustrated in Figure 4, have one characteristic in common; they each contain three areas of detail and it is the user's task in the exploration phase of the interview to locate these areas. Later the subjects were asked to make a movie showing the detail in context.

SIGNS. This environment consists of a three objects placed on a regular grid. These objects resemble roadsigns and each has an area of detail placed on one of its sides.

MAZE. This environment consists of a T shaped hallway. The three areas of detail are placed at different locations on the inside walls. From the user's starting position the entrance to the hallway appears as the center square in a checkerboard pattern. This checkerboard is opaque and only appears transparent in Figure 4 to reveal the T shaped hallway.

CUBE. This environment contains a single cube with details on three of its faces.

Evaluation Methodology

In order to evaluate the cognitive dimensions of the three interaction metaphors we decided to make use of a technique known as intensive "semi-structured" interviewing [2]. This technique, as its name suggests, involves interviewing users under controlled (structured)

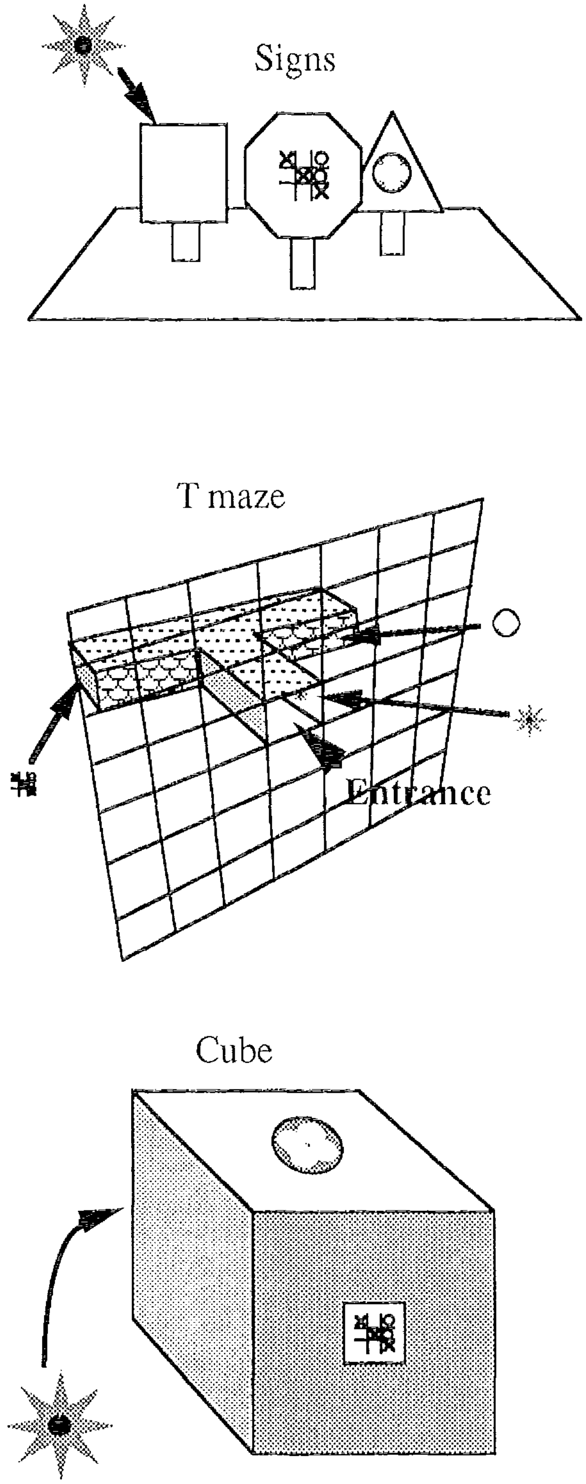

Figure 4. The three "toy worlds" are shown together with the locations of the detailing.

conditions, while at the same time providing some scope for the subject to provide creative input to the process.

The use of this evaluation methodology is a deliberate reaction against experimental studies which, because of the demands of experimental design, are constrained to ask very limited questions. The goal of intensive interviewing as an evaluation technique is to ask many meaningful questions of a few subjects, rather than asking a single (often trivial) question of many subjects.

A second difference in our approach from measurement based studies is that we consider individual differences to be important and analyzed them as such. Again this differs from the kind of study where a large number of subjects are distributed across experimental conditions, and measurements of responses are averaged across subjects within each condition (which makes an implicit assumption of homogeneity of the subject population). Because of this we do not require a large number of subjects, rather we require a small number each of whom we subject to an extended interactive interviewing session. Also, we do not try to hide individual differences, rather we try to gain knowledge from the variety of ways in which subjects respond to a particular situation. Because of this, in selecting subjects, we deliberately sought a variety of previous experience. We subject our data to *content analysis* [10] which essentially involves a careful semantic analysis of the actions and comments of the different subjects involved in different situations.

Subjects

The 7 subjects were chosen for the variety of their experiences. In particular, we wished to obtain subjects with and without computer experience, with and without mouse experience, with and without experience making movies and with and without experience flying aircraft or using flight simulators. We obtained information on the relevant experience of each subject by asking a series of questions and recording the answers. Space considerations prohibit the full reproduction of this information here. All but one of the subjects, were paid $30.00 for participation in the three hour session.

The Structure of the Interview

The product of the three interaction metaphors with the three toy environments yields nine distinct combinations. The interaction metaphors were evaluated by requiring subjects to explore each of the environments with each of the metaphors while making verbal comparisons and critical comments. The subject were given a brief introduction to a particular interaction metaphor, following which they were given the three scenes

	Environment in hand	Eyeball in hand	Flight vehicle control
Cube			
Control	+6	-2	-4
Movies	+4	0	-4
Exploring	+5	-2	-3
Maze			
Control	-6	0	+6
Movies	-7	+1	+6
Exploring	-6	0	+6
Signs			
Control	+2	-0.5	-1.5
Movies	-1	-1	+2
Exploring	0	0	0

Table 1

in succession. This was repeated with the two other metaphors and the order of presentation was randomized accross subjects. After each exploration session subjects were asked to make a movie showing each of the areas of detail and its relationship to the other areas of detail and to the rest of the scene. These interview sessions lasted approximately three hours and were videotaped in their entirety.

Results

In our preliminary analysis of the results we find no overall winner. There were also no real losers, each of the metaphors could be used within twenty minutes by users with little or no experience with computer systems or three-dimensional graphics.

The numbers given in Table 1 are the result of a set of questions which were asked at the end of the interview session. We discuss them first because they provide a kind of overview. Table 1 gives the compiled results obtained by asking each of the subjects about each scene in turn: how much they felt in control, how easy it was to make the movie, how easy it was to explore. Subjects were asked to choose the best and worst metaphors given each scene and these judgements were coded by scoring +1 when a metaphor was judged the best, a -1 when it was judged the worst, and a 0 otherwise.

One result which is apparent from Table 1 is the high correlation between the responses for the three questions. The Kendall coefficient of concordance gives this an agreement value of 0.88 which is highly significant ($p < 0.01$) This suggests that the ease of control, the ease of movie making and the ease of exploration have similar task constraints.

The other striking result is an interaction; "flight" was judged best for the MAZE but worst for the CUBE while "scene in hand" was judged best for the CUBE but worst for the MAZE. The reasons for this are discussed below where we present a condensed version of our detailed analysis of the videotape data. In the following analysis for each metaphor in turn, we discuss the constraints and affordances of our implementation, the cognitive properties of the metaphor, and the judged usefulness in making movies.

Eyeball in hand

Constraints and Affordances. Our implementation of the eyeball in hand metaphor imposes awkward physical contraints. It is necessary that the boundaries of the scene be located inside the boundaries of the device's domain, which in our case was a hemisphere with a radius of approximately two metres. This limitation prevents the viewpoint from being a great distance from any part of the scene. Also because this mapping relies on a one-to-one mapping of the hand to the virtual viewpoint it can be very hard to get a steady close look at a detail — this is because of human limitations in fine control of hand position and limitation in the resolution of the 3Space.

Another physical constraint is the movement of the user. To move the camera around a scene the user must physically walk around or reach around the virtual scene in the physical domain. This movement combined with complex hand movement, what one subject described as "physical contortions", is often necessary to get the desired view. Most complaints about the metaphor were related to this problem. The subjects preferred being able to sit down and explore the scene with a minimum of physical exertion and discomfort which they could do with both of the other metaphors.

Cognitive Properties. One of the strong points of this metaphor is also its weak point. The metaphor requires that the user imagine a virtual scene placed in the physical space in front of the monitor. There appears to be large individual differences in the ability of subjects to do this. The use of an actual physical model of the scene would eliminate this problem but this must be excluded for those cases where the scene to be explored is unknown. Besides which, the point of computer graphics is usually to avoid the need for physical models.

The most striking evidence that mental visualization of the scene plays a crucial role is that three of the subjects acted as if the scene were actually present. They made considerable efforts to keep out of the space inhabited by the virtual scene, and one subject, while exploring the SIGNS scene, would stretch uncomfortably rather than step forward and stand "in" the virtual scene.

The most serious cognitive problem with the eyeball in hand metaphor occurs when the 3Space is turned to face the user. This occurs when the user wants to view the back or bottom of a scene. The problem is that there is a disconcerting mismatch between hand motion and the visual feedback which results. In this situation the relative motion of the scene seems to be completely reversed and this is very disorienting. This seems to be because the actual viewpoint of the user, facing the monitor, is the reverse of the hand held virtual viewpoint.

There were variations in the mental imagery which was adopted by at least two of the subjects carrying out the exploration task using this metaphor. Sometimes these subjects tended to think of the device, not as an eyeball, but as a hand-held camera that they were walking through the maze, and there was a mismatch between the actual scale of the scene with respect to the bat and the perceived size of the camera which caused them some difficulty.

The eyeball metaphor is relatively easy to learn initially. There was no unanimous agreement by the subjects for the easiest metaphor to learn but the majority of the votes went to this one.

When asked which of the scenes this metaphor is most appropriate for, subjects seemed to be polarized into those who thought the MAZE was the most appropriate and the CUBE was the least appropriate, and those who felt the reverse was true.

Movie Making. When movie-making is considered, this metaphor was judged neither the best nor the worst. The direct correspondence between viewpoint motion and the user's movements makes it difficult to create smooth motions and a "jerky" movie is often the result. A smooth ease-in and ease-out of camera positions is virtually impossible.

Environment in hand

Constraints and Affordances. Our implementation of this metaphor only allows a single center of rotation for the entire scene. This is no problem with the cube scene where it seems natural to have the center of rotation at the center of the cube. However, with large or complex scenes, navigation is difficult when the viewpoint is far from the center of rotation. In particular, as the viewpoint is moved to a part of the scene which is distant from the center of rotation, the effects of rotating the bat are exaggerated, small angular changes result in significant positional changes and this tends to be very disorienting. One of our subjects actually gave up trying to navigate the maze with this metaphor. It is possible to imagine extensions to this metaphor by, for example, allowing the subject to place "handholds" on convenient surfaces. This would allow the subjects to get the feeling of directly manipulating the scene. We decided not to do this because it would require considerable computational power to implement and this would, we felt, drastically reduce the complexity of the scenes that could be manipulated in real time.

There is also a physical constraint involved. When trying to do a rotation of the scene through a large angle the user must rotate the bat completely around or ratchet. In either case the subjects found that the physical movement involved was often uncomfortable and was the major complaint concerning this metaphor. Additionally, ratchetting was found to be a tiring task.

Cognitive Properties. It appears that "scene in hand" is a natural metaphor for discrete hand sized objects. The CUBE is perceived by subjects to be approximately hand sized and six out of seven subjects judged this metaphor to be the best for working with this particular scene. Conversely it is considered be the worst for the MAZE environment. The metaphor received a mixed reaction from our subjects when used with the SIGNS scene. For some subjects a problem appeared to be caused by a mismatch between the "scene in hand" metaphor and the perceived size of the scene. The subjects felt that the SIGNS scene was very large and that it was therefore unnatural to be able to translate and rotate it so easily.

The subjects also expressed difficulty in rotating and translating at the same time. Because of the speed of response of the scene to the movements of the bat, simultaneous rotation and translation was found to be disorienting in the MAZE and the SIGNS. Subjects would attempt to separate translations from rotations and find this difficult to do.

Movie Making. This metaphor was rated the worst for movie-making. It was found to be especially hard to make a movie showing the detailing in the MAZE scene and the use of ratchetting produces "jerky" movement.

Flying vehicle control

Constraints and Affordances. This metaphor is unlike the others in that the subject controls the velocity of the viewpoint rather than it position and orientation. Subjects did not find this metaphor exceptionally difficult to learn although the concept was not as simple to pick up as the "eyeball in hand" metaphor. This metaphor was judged to be the best with the MAZE scene. The ability to move slowly and to easily make small adjustments were contributing factors to this concensus. The most disorienting aspect of the MAZE scene is passing through the walls of the corridors and this occurred far less frequently with "flight".

The use of the 3Space with this metaphor is less tiring than with "eye in hand" and "scene in hand" since it can be held in a comfortable position and does not require continuous movement by the user. On the other hand, this metaphor does require patience. With the "eyeball in hand" metaphor the user could move directly from one desired viewpoint to another very quickly if the user was capable of envisioning the virtual scene in the physical domain. The "flight vehicle control" metaphor does not allow this. The user can set the velocity of the viewpoint but must still wait until the viewpoint "flies" to the desired position.

One problem which the subjects experienced with the implementation was the lack of an alternative feedback on their velocities in the six dimensions. The only feedback was the visual flow feedback from the motion of the viewpoint. In a complex scene such as the MAZE this is not as severe a problem as it is in a simple enviroment such as the CUBE.

The majority of the subjects found this metaphor was the one least suited for the CUBE scene and this appeared to be due to the difficulty of flying around a single object in space while maintaining it in the line of sight. A complex interaction between angular and translational velocities is necessary to do this well. The best technique found by us and some of our subjects was to move sideways away from the object and rotate towards it at the same time. This combination of lateral and angular velocities results in an orbit about the object while keeping the viewpoint always toward the object itself. This technique is not an easy one to master in the limited amount of time the user's have to deal with each scene and the difficulty of learning it may have biased our results on the appropriateness of this metaphor with this scene. As an example of this, we noted that two subjects who explored the SIGNS environment immediately after the CUBE environment used the same orbitting technique on this environment and showed steady improvement in their skill with this task.

Cognitive Properties. An interesting cognitive phenomenon with "flight" is that subjects appear to adopt different metaphorical analogies depending on the scene. While exploring the SIGNS environment most subjects described the experience as being similar to flying. This led to some problems for our subjects with pilotting experience who had a difficult time realizing that they could fly backwards, sideways and up and down as well as remain stationary. One subject who has his pilot license even banked his viewpoint when he made a turn and rarely used any other type of movement other than forward movement. Another difficult concept for these subjects was flying toward the signs to examine the details and even flying through the signs rather than trying to avoid these obvious obstacles to the flight path. Interestingly, the subjects without pilotting experience had no difficulty with these concepts, thus actual flying experience may be a handicap more than an asset.

However our "flight vehicle control" interface did not always provide the perception of flight. While exploring the MAZE environment the majority of subjects describe the metaphor as being similar to walking down the corridors, instead of flying. This is obviously a cognitive consideration: "You walk down corridors, you don't fly down them". Additionally many subjects felt a cognitive constraint with the CUBE environment. The subjects felt that the cube was a small baseball-sized object and the concept of flying around such a small object seemed unnatural and some subjects would perceive the object as rotating, rather than themselves flying around it. The concept of actually rotating the object itself (as in the "environment in hand" metaphor) seemed much more natural.

A cognitive (and physical) advantage of this metaphor is the large amount of freedom that it affords to the user. The user can go anywhere in the scene at virtually any speed that is desired. A couple of our subjects described it as follows: "You believe you can do things compared to using the other metaphors." and "There is no feeling of restriction."

Movie Making. The fact "flying vehicle control" allows the user to control velocity, rather than position, and does so in a non-linear fashion, so that very slow speeds can be obtained, allows users to ease-in and ease-out the camera positions in a very smooth manner which is ideal for movie-making. This was the main reason cited for the general concensus that this metaphor is the best for this task. Our preliminary analysis of the movies made with each of the metaphors also indicates that this metaphor has the best results in terms of the quality of the movies produced.

Concluding Remarks

Perhaps the most interesting aspect of the results is the way in which metaphor appeared to influence the user's behaviour. The most striking examples of this were the subject who banked the viewpoint when turning corners when using the "flying vehicle control" metaphor, and the subject who carefully avoided placing his body in the same space as the virtual environment.

In considering the subjects' behaviour with each of the three scenes it seems as if they form a continuum in terms of the task demands. The SIGNS lies between the MAZE and the CUBE. The interface metaphors produced a strong interaction with this continuum, so that the "scene in hand" metaphor is strongly preferred with the CUBE, whereas the "flying vehicle control" metaphor is strongly to be preferred with the MAZE. This suggests that when designing interactions around a spatial metaphor considerable attention should be tied to cognitive conflicts which may result.

One final observation on the nature of metaphor. One of the criterion for a useful metaphor in a user interface is whether it can be extended to meet new requirements. A good metaphor is thus one which has many useful elabourations. An application which we have for the present work is to extend our exploration interface into a tool which allows the editing of data together with its exploration. Our immediate problem is to delete outlying data points produced by high volume ocean mapping systems. We wish to extend the metaphor from exploration to exploration combined with manipulation. The flight metaphor has an obvious extension — especially for those who play video games — into shooting down the bad data. In an alternative extension of the metaphor we can imagine picking up data objects and flying them to new locations. The scene in hand metaphor can also be extended. If we can detach our linkage to the scene, we can move a cursor within the scene to a data object and once selected the object can be delete or moved (this implies a scene which is stored as a two level heirarchy). The eyeball in hand metaphor is unique in that it offers no natural extension to data manipulation.

Acknowledgements

We wish to acknowledge the assistance of Roy Nicholl, for software development, and help with the interview process. Mike MacDonald provided invaluable systems support and Dr. John M. Kennedy, of the University of Toronto, provided insights into the nature of metaphor. Leonard Slipp helped us to clarify the manuscript and to sort out some ideas on the task constraints. Funding was provided by the Natural Sciences and Engineering Research Council of Canada.

References

[1] Badler, N.I., Manoochehri, K.H., and Baraff, D., Multi-Dimensional Input Techniques and Articulated Figure Positioning by Multiple Constraints. *Proc. 1986 Workshop in Interactive 3D Graphics*, F. Crow and S.M. Pizar, eds, ACM, New York, pp 151–169, 1986.

[2] Brenner, M., Intensive Interviewing, In *The Research Interview: Uses and Approaches*, M. Brenner, J. Brown and D. Canter, Eds. Academic Press: London, pp 147–162, 1985.

[3] Brooks, F.P., Walkthrough — A Dynamic Graphics System for Simulating Virtual Buildings. *Proc. 1986 Workshop on Interactive 3D Graphics*, F. Crow and S.M. Pizar, eds, ACM, New York, pp 9–21, 1986.

[4] Brooks, F.P., Grasping Reality Through Illusion: Interactive Graphics Serving Science. *CHI'88 Proceedings*, Special Issue of the SIGCHI Bulletin, pp 1–11, 1988.

[5] Chen, M., Mountford, S.J., and Sellen, A., A study of Interactive 3D Rotation Using 2D Control Devices. *Computer Graphics*, 22(4), pp 121–129, 1988.

[6] Evans, K.B., Tanner, P.P., and Wein, M., Human Spatial Orientation. Tablet based valuators that provide one, two, or three degrees of freedom. *Computer Graphics*, 15(3), pp 91–97, 1981.

[7] Fisher, S.S., McGreevy, M., Humphries, J., and Robinett, W., Virtual Environment Display System. *Proc. 1986 Workshop on Interactive 3D Graphics*, F. Crow and S.M. Pizar, eds, ACM, New York, pp 77–88, 1986.

[8] Kelley, C.R., *Manual and Automatic Control*. Wiley, New York, 1968.

[9] Manktelow, K., and Jones, J., Principles from the Psychology of Thinking and Mental Models. In *Applying Cognitive Psychology to User Interface Design*, M.M. Gardiner and B. Christie, eds., Wiley, New York, 1987.

[10] Mostyn, B., The Content Analysis of Qualitative Research Data: A Dynamic Approach. In *The Research Interview: Uses and Approaches* J.M. Brenner, T. Brown and D. Canter, eds, Academic Press: London, pp 115–145, 1985.

[11] Schmandt, C., Spatial input display correspondence in a stereoscopic computer graphic workstation. *Computer Graphics*, 17(3), pp 253–259, 1983.

[12] Sutherland, I.E., Head Mounted Three Dimensional Display. *Proc. of the Fall Joint Computer Conference*, Vol. 33, pp 757–764, 1968.

[13] Veldhuyzen, W., and Stassen, H.G., The internal model concept: An application to modelling human control of large ships. *Human Factors*, 19, pp 367–380, 1977.

[14] Ware, C., Using Hand Position for Virtual Object Placement. (A report of empirical results using the Bat for placement). Submitted for publication.

[15] Ware, C. and Jessome, D.R., Using the Bat: A six-Dimensional mouse for Object Placement, *IEEE Computer Graphics and Applications*, 8,6, pp 155–160, 1988.

[16] Zimmerman, T.G., Lanier, J., Blanchard, C., and Harvill, Y., A Hand Gesture Interface Device, *CHI+GI Proceedings*, Toronto, pp 189–192.