

The Influence of Dynamic Shadows on Presence in Immersive Virtual Environments

Mel Slater, Martin Usoh, Yiorgos Chrysanthou¹,
Department of Computer Science, and
London Parallel Applications Centre,
QMW University of London,
Mile End Road,
London E1 4NS, UK.

Abstract. This paper describes an experiment where the effect of dynamic shadows in an immersive virtual environment is measured with respect to spatial perception and presence. Eight subjects were given tasks to do in a virtual environment. Each subject carried out five experimental trials, and the extent of dynamic shadow phenomena varied between the trials. Two measurements of presence were used - a subjective one based on a questionnaire, and a more objective behavioural measure. The experiment was inconclusive with respect to the effect of shadows on depth perception. However, the experiment suggests that for visually dominant subjects, the greater the extent of shadow phenomena in the virtual environment, the greater the sense of presence.

Keywords. Virtual environment, virtual reality, presence, shadows, depth perception.

1. Introduction

We describe an experiment to examine the effect of shadows on two different aspects of the experience of immersion in a virtual environment (VE): depth perception and presence. It is well-known that shadows can significantly enhance depth perception in everyday reality [1,5,7]. Shadows provide alternative views of objects, and provide direct information about their spatial relationships with surrounding surfaces. VR systems typically do not support shadows, and yet potential applications, especially in the training sphere, will require participants to make judgements about such relationships. Even the simple task of moving to an object and picking it up can be problematic when observers cannot easily determine their own distance from the object, or its distance from surrounding objects. We

¹ emails: mel, bigfoot, yiorgos @ dcs.qmw.ac.uk. URL: <http://www.dcs.qmw.ac.uk/>

introduce dynamic shadows to examine whether such task performance can be enhanced.

We have argued elsewhere [8] that presence is the key to the science of immersive virtual environments (virtual reality). We distinguish, however, between immersion and presence. Immersion includes the extent to which the computer displays are extensive, surrounding, inclusive, vivid and matching. The displays are more extensive the more sensory systems that they accommodate. They are surrounding to the extent that information can arrive at the person's sense organs from any (virtual) direction. They are inclusive to the extent that all external sensory data (from physical reality) is shut out. Their vividness is a function of the variety and richness of the sensory information they can generate [11]. In the context of visual displays, for example, colour displays are more vivid than monochrome, and displays depicting shadows are more vivid than those that do not. Vividness is concerned with the richness, information content, resolution and quality of the displays. Finally, immersion requires that there is match between the participant's proprioceptive feedback about body movements, and the information generated on the displays. A turn of the head should result in a corresponding change to the visual display, and, for example, to the auditory displays so that sound direction is invariant to the orientation of the head. Matching requires body tracking, at least head tracking, but generally the greater the degree of body mapping, the greater the extent to which the movements of the body can be accurately reproduced.

Immersion also requires a self-representation in the VE - a Virtual Body (VB). The VB is both part of the perceived environment, and represents the being that is doing the perceiving. Perception in the VE is centred on the position in virtual space of the VB - e.g., visual perception from the viewpoint of the eyes in the head of the VB.

Immersion is an objective description of what any particular system does provide. Presence is a state of consciousness, the (psychological) sense of being in the virtual environment. Participants who are highly present should experience the VE as more the engaging reality than the surrounding world, and consider the environment specified by the displays as places visited rather than as images seen. Behaviours in the VE should be consistent with behaviours that would have occurred in everyday reality in similar circumstances.

Presence requires that the participant identify with the VB - that its movements are his/her movements, and that the VB comes to "be" the body of that person in the VE. We speculate that the additional information provided by shadows about the movements of the VB in relationship to the surfaces of the VE can enhance this degree of association, and hence the degree of presence. However, we were unable to test this in the current experiment. We do, however, consider the proposition that shadows, increasing the degree of vividness of the visual displays, will enhance the sense of presence.

2. Experiment

2.1 Scenario

The experimental scenario consisted of a virtual room, the elevation of which is shown in Figure 1. Five red spears are near a wall, but behind a small screen. Another green spear is at position G. The subject begins the experiment by moving to the red square (X), and facing the spears. The instruction is to choose the spear nearest the wall, observing from position X. Having chosen that spear, the subject moves towards it, picks it up and returns to X. There the subject turns to the left, facing a target on the far wall. The subject must orient the spear to point approximately towards the target, fire and guide it towards the target by hand movements. The instructions were that the spear must be shot at the target, and that it must be stopped the instant that its point hit the target. Finally, the subject must bring the green spear to position X. This was repeated six times for each subject.

Prior to the start of the experiment each subject was given a sheet explaining these procedures, and the first run was for practice, the experimenter talking the subject through the entire scenario. Runs 1 through 5 were carried out by the subject without intervention by the experimenter. Between each run the subject was advised to relax with closed eyes, either with or without the head-mounted display (HMD, see below), although all but one continued to wear it during the two minutes that it took to load the program for the subsequent run. Each of the five runs were the same apart from the distances of the red spears from the wall. Also, some runs displayed dynamic shadows of the spears and the small screen, while others did not.

Eight subjects were selected by the experimenters asking people throughout the QMW campus (in canteens, bars, laboratories, offices) whether they wished to take part in a study of "virtual reality". People from our own Department were not included.

Table 1
Runs of the Experiment for Each Subject
1,2,3,4 denotes the four point-light positions of Figure 1
0 denotes no shadows

Subject	No. shadow scenes / 5	Run 1	Run 2	Run 3	Run 4	Run 5
1	1	0	0	2	0	0
2	1	0	0	2	0	0
3	2	0	0	2	3	0
4	2	0	0	2	3	0
5	3	0	1	2	3	0
6	3	0	1	2	3	0
7	4	0	1	2	3	4
8	4	0	1	2	3	4

The design is shown in Table 1, which indicates the positions of the point-light source for those runs that included shadows. Note that of the 40 runs, 20 included shadows.

2.2 Spatial Variables and Hypotheses

The variables measured in order to assess the effects of shadows on spatial judgement were as follows:

Spear Selected. S: the spear selected from observation position X. The spears ranged from 50 cm to 90 cm from the wall, positioned with 10 cm variations. The small screen in front of the spears obscured the positions where they touched the floor, for any subject standing at position X. Also, because their distances from the wall varied only slightly, their heights, as judged from position X would look the same. It was therefore very difficult to judge which spear was nearest the wall. Variable S was the rank order of the spear chosen, where 1 would be the nearest to the wall, and 5 the furthest.

The hypothesis was that subjects would be able to use the shadows of spears on the walls to aid their judgement about the closeness to the walls, so that those runs that included shadows would result in a greater number of correct spears being chosen.

Distances from Target. C: this is the distance of the point of the spear from the centre of the target at the position that it was stopped in flight by the subject.

The hypothesis was that the subjects would be able to use the shadow of the spear in flight, especially its shadow on the target wall, to help guide the spear towards the target. Therefore, the mean distance should be less for the shadow runs than for the non-shadow runs.

D: this is the distance that the point of the spear was behind or in front of the target at the position that it was stopped by the subject.

The hypothesis is as for C, except that here we would expect a greater shadow effect since the action required to stop the spear in flight (releasing a button on the hand-held 3D mouse) is simpler than that involved in guiding the spear to the bulls eye. Moreover, at the moment the spear point touched the target wall, it would also meet its shadow.

2.3 Presence Variables and Hypotheses

In previous studies we have used subjective reported levels of "presence" based on a questionnaire. In this method subjective presence was assessed in three ways: the

sense of "being there" in the VE, the extent to which there were times that the virtual world seemed more the presenting reality than the real world, and the sense of visiting somewhere rather than just seeing images. In the present study these three basic determinants were elaborated into six questions, each measured on a 7-point scale, where lowest presence is 1, and highest is 7 (see Appendix A). The overall presence score (P) was conservatively taken as the number of high (6 or 7) ratings amongst the six questions, so that $0 \leq P \leq 6$.

Although we have obtained good results with such subjective measures before, in the shadow experiment we introduced in addition a more "objective" measurement of presence. This was achieved by having one particular object (a radio) in both the real world of the laboratory in which the experiment took place and the virtual world of the room with spears.

Just before the practice run the subjects were shown a radio on the floor against a large screen in the laboratory. They were told that they would see "the radio" in the virtual world, and that occasionally it would switch itself on. Whenever they heard the sound they should point towards "the radio", and press a button on the hand-held mouse. This would act as an "infra-red" device to switch the radio off. Before they entered into the VE the radio was momentarily switched on, deliberately not tuned to any particular channel therefore causing it to play an audible but meaningless tone. Each time that the subject entered into the VE, i.e., at the start of each run they were told: "Orient yourself by looking for the red square on the floor and the radio". The radio was placed in the VE at the same position relative to the red square as the real radio was to the position of the subject just before entering the VE.

At four moments during the experiment, always while the subject was (virtually) on the red square, the real radio was moved to one of four different positions. These were 1m apart from each other, on a line coincident (in the real world) with the small screen by which the radio was located (in the virtual world). The ordering was selected randomly before the start of the experiment. The virtual radio was always in the same place. Therefore the subject would hear the sound coming from a different location compared to the visible position of the radio. The idea is that (other things being equal), a high degree of presence would lead to the subject pointing towards the virtual radio rather than the real one. Hence we tried to cause and use the conflict between virtual and real information as an assessment of presence. Those (two) subjects who did ask about the contradiction were told "Just point at where you think the radio is". Throughout, both the real radio and the virtual radio were referred to as "the radio", deliberately allowing for a confusion in the minds of the subjects.

It is important to note that we mean "presence" in a strong behavioural sense with respect to this measurement. The questionnaire attempts to elicit the subject's state of mind. The radio method though is concerned only with their behaviour. If they pointed to the virtual radio because of a need to obey the experimenter, or because it

was a matter of "playing the game", then so be it. Provided that they act in accordance with the conditions of the VE, this is behavioural presence.

Let R be the angle between the subject's real pointing direction and the direction to the real radio. Let V be the angle between the subject's virtual pointing direction and the direction to the virtual radio. Small V therefore occurs when the subject points towards the virtual radio. We use $P_a = R/V$ as the measurement of the extent to which the subject tends towards the virtual radio - a small V in comparison to R would result in large P_a . Therefore larger values of P_a indicate greater tendency towards the virtual.

There were two hypotheses relating to P_a : First, that it would correlate positively with P , and second that the greater exposure of the subject to shadows, the greater the value of P_a . Of course, we would also expect that the greater the exposure to shadows, the greater the value of P .

2.4 Representation System Dominance

A clear objection to this procedure is that it could be measuring the extent of visual or auditory dominance rather than presence. Faced with conflicting information from two senses, the resulting action is likely to depend on which sensory system is "dominant". In previous work [9,10] we have explored the relationship between dominant representation systems and the extent of subjective presence, and have always found a very strong relationship. This is based on the idea that people differ in the extent to which they require visual, auditory or kinesthetic/tactile information in order to construct their world models, and that each person may have a general tendency to prefer one type of representation (say visual) over another (say auditory). We found that in experiments where the virtual reality system presented almost exclusively visual information, the greater the degree of visual dominance the higher the sense of presence, whereas the greater degree of auditory dominance, the lower the sense of presence.

In this shadow experiment therefore we employed an updated version of the questionnaire we used in [10] which is given to the subjects before attending the experimental session. This questionnaire attempts to elicit their preferences regarding visual, auditory and kinesthetic modes of thinking. It presents 10 situations, each one having three responses (one visual, one auditory, and one kinesthetic response). Subjects are asked to rank their most likely response as 1, next most likely as 2, and least likely as 3. From this a V score is constructed as the total number of $V=1$ scores out of 10, and similarly for A and K . Alternatively the sums of the responses may be used. These V and A variables can therefore be used to statistically factor out the possible influence of visual or auditory dominance on the radio angles.

The hypothesis with respect to V, A and K would be that V and K would be positively correlated with presence (however it is measured) whereas A would be negatively correlated, in line with our previous findings. Note that by construction, there are only 2 degrees of freedom amongst V, A and K.

3. Apparatus

3.1 Equipment

The experiments described in this paper were implemented on a DIVISION ProVision system, a parallel architecture for implementing virtual environments running under the dVS (v0.3) operating environment. The ProVision system is based on a distributed memory architecture in which a number of autonomous processing modules are dedicated to a part of the virtual environment simulation. These processing modules or Transputer Modules (TRAMs) are small self-contained parallel processing building blocks complete with their own local memory and contain at least one Inmos Transputer which may control other specialised peripheral hardware such as digital to analog converters (DAC). Several modules exist. These include:

- the module to act as the module manager.
- the DAC module for audio output.
- polygon modules for z-buffering and Gouraud shading.
- application specific modules for the user applications.

The dVS operating environment (Grimsdale, 1991) is based on distributed Client/Server principles. Each TRAM or processing cluster is controlled by an independent parallel process known as an Actor. Each provides a set of services relating to the elements of the environment which it oversees. Such elements presently consist of lights, objects, cameras, controls (i.e. input devices), and collisions between objects. Thus, an Actor provides a service such as scene rendering (visualisation actor). Another Actor may be responsible for determining when objects have collided (collision actor) and yet another for hand tracking and input device scanning. All these Actors are co-ordinated by a special Actor called the Director. Communication between the different Actors can only be made via the Director. The Director also ensures consistency in the environment by maintaining elements of the environment which are shared by the different Actors.

The ProVision system includes a DIVISION 3D mouse, and a Virtual Research Flight Helmet as the head mounted display (HMD). Polhemus sensors are used for position tracking of the head and the mouse. The displays are colour LCDs with a 360×240 resolution and the HMD provides a horizontal field of view of about 75 degrees.

All subjects saw a VB as self representation. They would see a representation of their right hand, and their thumb and first finger activation of the 3D pointer buttons would be reflected in movements of their corresponding virtual finger and thumb. An example is shown in Plates 1 and 2. The hand was attached to an arm, that could be bent and twisted in response to similar movements of the real arm and wrist. The arm was connected to an entire but simple block-like body representation, complete with legs and left arm. Forward movement was accompanied by walking motions of the virtual legs. If the subjects turned their real head around by more than 60 degrees, then the virtual body would be reoriented accordingly. So for example, if they turned their real body around and then looked down at their virtual feet, their orientation would line up with their real body. However, turning only the head around by more than 60 degrees and looking down (an infrequent occurrence), would result in the real body being out of alignment with the virtual body.

The 3D mouse is shaped something like a gun. There is a button in the position of the hammer, which is depressed by the thumb. This causes forward motion in the direction of pointing. There is a button on each side of this central thumb button, each activated by the thumb. The left one was used to fire the spears - while this button was depressed the spear would move in a direction determined by hand orientation. The spear would stop on release of this button, and could not be activated again, thus giving the subject one chance per spear. The right thumb button was used as the "infra-red" radio switch. Corresponding to the trigger is a button for the forefinger. This is used to pick objects - squeezing this finger button while the virtual hand intersects an object results in the object attaching to the hand. Subjects were able to master these controls very quickly.

3.2 Shadow Algorithm and Frame Rates

The shadow algorithm is described in detail elsewhere [3]. It is based on a dynamic Shadow Volume BSP tree [2], constructed from polygons in arbitrary order, that is without the necessity of a separate scene BSP tree. Shadows are created as polygons in object space. Creation of new shadows and changes to shadows are communicated dynamically to the renderer via the Director.

For reasons described below, the entire scene was small, consisting of 413 triangles, of which only 52 would be likely to influence shadow creation. The frame rate achieved without shadows was 9Hz. The frame rate with shadows, 6 to 8Hz, was not very satisfactory, but due to the particular version of the dVS software architecture in use on this machine at the time of the experiment.

Without rendering the shadow algorithm runs on this machine at a frequency of between 19 and 21Hz depending on the complexity of the view at any moment. The renderer does not however run at this frequency during dynamic changes of a virtual object, due to update problems associated with the extant implementation of the dVS dynamic geometry object. Therefore, when rendering and the associated

communication time is included, the frame rate is 6 to 8Hz. (A new version of dVS is intended to solve this problem).

dVS v0.3 maintains the concept of a "dynamic geometry object". This is a vertex-face structure representing a (possibly empty) set of polygons. The actual polygons belonging to this object can be created or modified at run time. When such a change is made to a dynamic object, there is an "update" generated that sends the object to the Director for distribution to the Visualisation Actor and then onto to the renderer.

Upon any change of a virtual object the shadow algorithm recomputes the shadow scene outputting any modified shadow polygons, i.e. any polygons that have been deleted and any that have been created. This information is transmitted to the shadow generation module which will mark deleted polygons as invisible to be re-used later by new shadow polygons. The module uses a linked list structure of dynamic objects - the shadow object. Each element in the list is a dynamic object consisting of 32 shadow polygons. This linked list structure is necessary in order to break down the entire list of potential shadow polygons into smaller chunks, rather than have one dynamic geometry object for all possible shadows, since the dynamic geometry implementation can only send updates of an entire dynamic object to the Visualisation Actor. Note that a change in one single shadow polygon will result in the communication of a complete 32-polygon dynamic object. If, unfortunately, 33 shadow polygons change, then two dynamic objects consisting of 64 polygons are communicated, and so on.

There is one important implication of this for the spatial judgement component of the experiment - obviously the spear travels more slowly when there are shadows. Without shadows the mean velocity is 92 cm/sec, and with shadows 47 cm/sec. Therefore it can be argued that differences in targeting performance might result from the velocity rather than the use of shadows. However, the effect of this can be examined statistically. With regard to the influence on presence we would argue that the slower frame rate in the case of shadows would tend to have a negative effect on presence.

4. Results

4.1 Spatial Variables

Spear Selected. Shadows made no difference at all to the selection of the "correct" spear (the one closest to the wall).

Distances from Target. Consider first C the distance of the point of the spear from the centre of the target. A regression analysis was used to examine the effect of velocity, showing that velocity within each of the shadow/ no-shadow groups was did not have a statistically significant effect. The mean distance without shadows is

152cm and 115cm with shadows. However, the difference between these two is not statistically significant.

Consider next D, the perpendicular distance of the point of the spear from the wall of the target. This could be positive (spear stops in front of the target) or negative, the spear stops behind). Carrying out a within-group regression analysis to examine the effect of velocity again shows that velocity is not statistically significant. The means are -39.9cm without shadows, and 3.3cm with shadows. The standard errors are 3.6 and 3.5 respectively and the difference is significant at 5%. The medians of the shadow and non-shadow D values are -3cm and -38cm respectively.

Although the within-group velocity appeared not to be statistically significant in each case, there is still some doubt about whether the inference about better performance in the case of shadows is safe. The variation of velocity within groups was not very great (the minimum and maximum velocities were 81.6 to 99.0 for the non-shadow group, and 36.0 to 60.4 for the shadow group). Subsequent experiments should attempt to produce a greater similarity in performance between the two groups.

4.2 Presence

Subjective Presence. P is the number of "high" questionnaire scores, as a count out of 6. We therefore treated P as a binomially distributed dependent variable, and used logistic regression.

In logistic regression [4], the dependent variable is binomially distributed, with expected value related by the logistic function to a linear predictor. Let the independent and explanatory variables be denoted by x_1, x_2, \dots, x_k . Then the linear predictor is an expression of the form:

$$\eta_i = \beta_0 + \sum_{j=1}^k \beta_j x_{ij} \quad (i = 1, 2, \dots, N) \dots\dots\dots(1)$$

where N (=8) is the number of observations. The logistic regression model links the expected value $E(P_i)$ to the linear predictor as:

$$E(P_i) = \frac{n}{1 + \exp(-\eta_i)} \dots\dots\dots(2)$$

where n (=6) is the number of binomial trials per observation.

Maximum likelihood estimation is used to obtain estimates of the β coefficients. The deviance (minus twice the log-likelihood ratio of two models) may be used as a goodness of fit significance test, comparing the null model ($\beta_j = 0, j = 1, \dots, k$) with any given model. The change in deviance for adding or deleting groups of variables may also be used to test for their significance. The (change in) deviance has an approximate χ^2 distribution with degrees of freedom dependent on the number of parameters (added or deleted).

Table 2

Logistic Regression Equations

$\hat{\eta}$ = fitted values for the presence scale

A = Auditory Sum, NS = number of shadows

Standard Errors shown in brackets

Model
$\hat{\eta} = 15.0 + 0.7*NS - 9.5*A$ <div style="display: flex; justify-content: space-around; width: 100%;"> (3.7) (0.4) </div>

Overall Deviance = 3.454, d.f. = 5

χ^2 at 5% on 10 d.f. = 11.070

Deletion of Model Term	Change in Deviance	Change in d.f.	χ^2 at 5% level
NS	4.123	1	3.841
A	9.088	1	3.841

Table 3*Normal Regression Equations* \hat{P}_a = fitted values for the angular discrepancy

NS = number of shadows

Group	Model
Visually dominant	$\hat{P}_a = -13.6 + 10.6*NS$ (3.7)
Auditory dominant	$\hat{P}_a = 9.427 + 0.08*NS$ (3.7)

Multiple Correlation Coefficient, $R^2 = 0.29$, d.f. = 36

Table 2 shows the result of the fit with P as the dependent variable, and the number of shadow runs (NS) and the auditory sum score (A) as the explanatory variables, across the 8 subjects. These were the only statistically significant variables found, and this supports the hypothesis that subjective presence is positively related with the shadow effect. As we have found previously, given this exclusively visual VE, the greater auditory dominance, as measured by the sum of A responses to the pre-questionnaire, the less the reported subjective presence.

Angular Discrepancy. Here we take P_a as the dependent variable and carry out a Normal regression with number of shadows (NS) and the representation system scores as the explanatory variables. NS proved once again to be significant and positively related to P_a . However, the V, A and K variables were not significant. Nevertheless it seemed important to try to rule out the possibility that the result with the angular discrepancy was simply due to visual or auditory dominance. Therefore a new factor was constructed, "sensory dominance" which has the value 1 if $V > A$ otherwise 2. Hence this directly refers to visual or auditory dominance. The result of the regression analysis including this was interesting: for those who were visually dominant, there is a significant positive relationship between P_a and NS, whereas there is no significant relationship for those who were dominant on the auditory score. This is shown in Table 3. (It so happened that 4 of the subjects were visually dominant).

5. Conclusions

There are three main issues : First, the point of this paper is not that we have an algorithm that can generate shadow umbrae rapidly in dynamically changing scenes. Even in this very small scene the rendering frame rate was no where near adequate on this particular architecture, though its performance is excellent on standard

workstations running under X11[3]. There is clearly a lot of work to do in the location of this algorithm in the dVS system architecture, in order to obtain maximum performance by minimising communication bottlenecks.

Second, although we have considered depth and spatial perception problems in the experiment, again, this is not the major point. It is more or less obvious, from everyday reality, and from perceptual studies that shadows do indeed enhance depth perception. Moreover, our experimental design in this regard was not ideal, since we did not control a factor (velocity) that potentially has an impact on the results.

Third, the real point of the experiment was the examination of the relationship between dynamic shadows and the sense of presence. This result is not obvious, and was motivated by the idea that presence is (amongst other things) a function of immersion, and immersion requires vividness. We used two independent measures - one subjective from the post-experiment questionnaire, and the other objective, as a ratio of angles of real to virtual pointing directions. Each method gave similar results, and the two measures were significantly correlated. Moreover, we found that for those people who were more visually dominant their (angular ratio) presence increased with exposure to shadows but that this did not hold for those who were dominant on the auditory scale. Increase in the subjective presence scale was also associated with an increase in shadow exposure, but with a decrease in the auditory scale. These results also support our earlier findings regarding the importance of the sensory system preferences in explaining presence.

We suspect that much stronger results on presence would have been obtained had we been able to allow the virtual body to cast shadows. However, this was not practical given the communication bottleneck problems discussed in §3.2.

If an application does not require presence, there is little point in using a virtual reality system. If a virtual reality system is used for an application, then there is little point to this unless it can be shown that a sense of presence is induced for most of the potential participants. Should the results of our shadow experiment be confirmed by later studies then it will have been shown that the great computational expense of shadow generation is worth-while for those applications where the participants are likely to be "visually dominant".

Acknowledgements

This work is funded by the U.K. Engineering and Physical Sciences Research Council (EPSRC), the Department of Trade and Industry, and DIVISION Ltd, through grant CTA/2 of the London Parallel Applications Centre. Thanks to David Sweeting of Aeronautical Engineering Department of QMW, for helping to find subjects for the experiment.

References

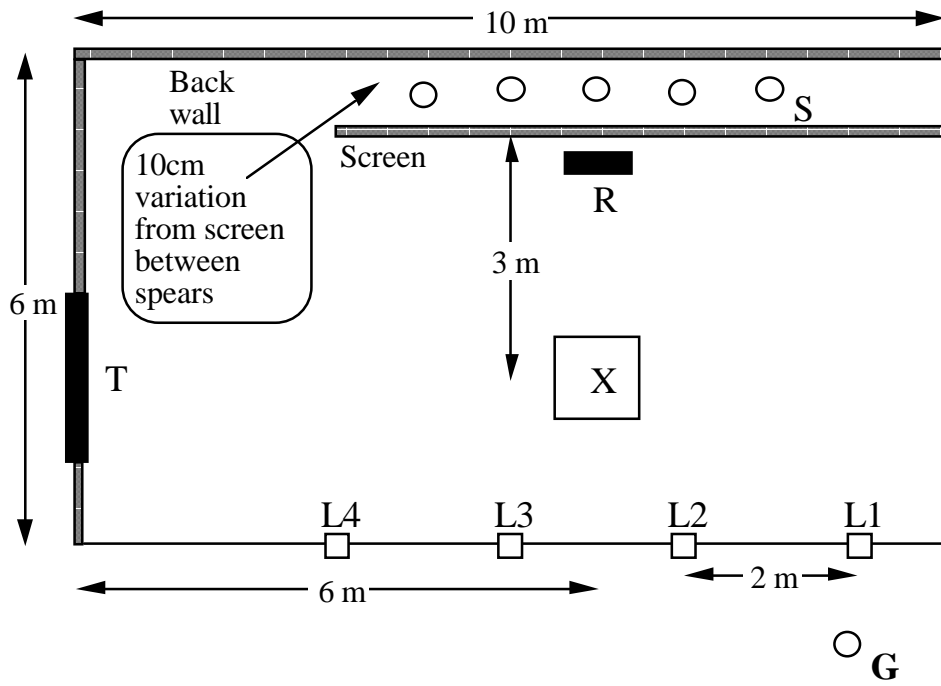
1. Cavanagh, P., Leclerc, Yvan G. (1989) Shape from shadows. *Journal of Experimental Psychology: Human Perception & Performance*, 15 (1), 3-27.
2. Chin, N., Feiner, S. (1989) Near Real-Time Shadow Generation Using BSP Trees, *Computer Graphics* 23(3), 99-106.
3. Chrysanthou, Y., Slater, M. (1995) Shadow Volume BSP Trees for Computation of Shadows in Dynamic Scene, *ACM SIGGRAPH Symposium on Interactive 3D Graphics* (April, 1995).
4. Cox, D.R. 1970. *Analysis of Binary Data*, London: Methuen.
5. Gregory, R.L. (1990) *Eye and Brain: The Psychology of Seeing*, Fourth Edition, Weidenfield and Nicholson, 182-187.
6. Grimsdale, C. (1991). dVS - Distributed Virtual environment System, *Proceedings of Computer Graphics 91 Conference*, London.
7. Puerta, A.M. (1989) The power of shadows: shadow stereopsis. *Journal of the Optical Society of America*, 6, 309 - 311.
8. Slater, M. A. Steed and M. Usoh (1994) Steps and Ladders in Virtual Reality, *ACM Virtual Reality Science and Technology (VRST)*, eds G. Singh and D. Thalmann, World Scientific, 45-54.
9. Slater, M. and M. Usoh, Representation Systems, Perceptual Position and Presence in Virtual Environments, *Presence: Teleoperators and Virtual Environments*, 2.3 MIT Press, 221-234.
10. Slater, M., M. Usoh, A. Steed, Depth of Presence in Virtual Environments, *Presence: Teleoperators and Virtual Environments*, 3.2 (in press).
11. Steuer, J. (1992) Defining Virtual Reality: Dimensions Determining Telepresence, *Journal of Communication* 42(4), 73-93.

Appendix A: Presence Questions

All questions were answered on a 1 to 7 scale, not reproduced here for space reasons.

1. Please rate *your sense of being there* in the virtual reality.

2. To what extent were there times during the experience when the virtual reality became the "reality" for you, and you almost forgot about the "real world" of the laboratory in which the whole experience was really taking place?
3. When you think back about your experience, do you think of the virtual reality more as *images that you saw*, or more as *somewhere that you visited* ?
4. During the course of the experience, which was strongest on the whole, your sense of being in the virtual reality, or of being in the real world of the laboratory?
5. When you think about the virtual reality, to what extent is the way that you are thinking about this similar to the way that you are thinking about the various places that you've been today?
6. During the course of the virtual reality experience, did you often think to yourself that you were actually just standing in a laboratory wearing a helmet, or did the virtual reality overwhelm you?



X - red square (judgement point)
 S - spear
 T - target
 G - green spear

* **E**
 E - scene entry point
 L - light (positions 1, 2, 3, 4)
 R - virtual radio

Figure 1
 Plan View of the Virtual Environment

Plates 1 and 2

Views of the virtual right arm as a box with shadows is approached.