What do you think you're doing? - Measuring perception in Fish Tank Virtual Reality

Michael Treadgold[1], *Kevin Novins*[1], *Geoff Wyvill*[1], *and Brian Niven*[2]

[1]Dept. Computer Science,
[2]Dept. Mathematics and Statistics,
University of Otago, Dunedin, New Zealand
miket5@cs.otago.ac.nz

Abstract

Fish Tank Virtual Reality (FTVR) systems use eye tracking to present perspective projections to a user that are tailored to his or her view position. These perspective views can create the illusion of a world behind the display screen – a virtual fish tank. However, the illusion is not perfect – users do not feel that the virtual objects are real. Despite this, users clearly are getting some indication of depth. What exactly are they perceiving?

In this paper, we review previous work in FTVR and identify key differences between the virtual world and the real world. We discuss the results of two previously published perceptual studies and a third of our own. We show that we still have an incomplete understanding of perception in Fish Tank Virtual Reality and that some important questions may be impossible to answer. Although FTVR clearly is useful, it should be used with caution if the task requires accurate and immediate 3D perception.

Keywords: *Fish Tank Virtual Reality, perception.*

1 Introduction

Fish Tank Virtual Reality (FTVR) systems are a class of virtual reality systems that track the position of the viewer's eyes and draw the correct perspective view of the virtual scene for that position. Figure 1 shows the vertex of a virtual object being projected onto the screen. It also shows that the projection will change if the user moves. This can create the illusion that the virtual object is located in the same real space as the user.

There are many advantages of FTVR systems over the more immersive type of virtual reality that requires a head mounted display (HMD):

- Versatility and Price - FTVR systems can be used on standard desktop computers, which makes integrating VR into the home and office simple.
- High Resolution - FTVR systems can provide more pixels per degree in the user's field of view than is possible in HMD systems. In addition, the light intensity per degree can be higher [12].

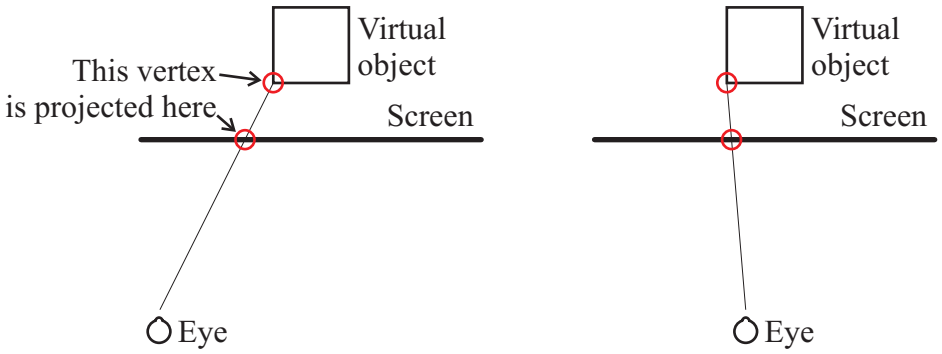

Figure 1: A virtual object being projected onto the screen from two different viewpoints (viewed from above).

- Minimal Motion Sickness - A high proportion of users of immersive VR experience some sort of motion sickness [14]. This is most likely caused by sensory conflict, but display update lag, low resolution, and full visual immersion may also play a part. In contrast, the most immersive FTVR system, the CAVE, produces motion sickness in only a very small proportion of users [1].
- Mobility and Communication - For many applications immersive VR is too involving. Since the user is tethered, his or her mobility in the real world is impaired. Having to don a glove and helmet make it difficult to switch between the virtual world and the real world [15]. As a result, communication between VR participants and non-participants is also hampered.
- Rotation Errors - In FTVR systems the projection plane, the screen, does not move with the viewer's position or orientation, so head or eye rotation creates no error. In HMD systems, a rotational tracking error produces a rotational error of the same magnitude in the image [1].

FTVR has been applied to 3D CAD systems such as JDCAD [10], large document retrieval systems such as Document Explorer [7], and a prototype Smart Kiosk [21]. Applications have also been suggested in the area of advertising [11]. Large FTVR systems, such as the CAVE [2], the Wedge [8], and the Responsive Workbench [9] have primarily been applied to scientific visualisation. The Cubby system [5] has been proposed as a candidate for medical visualisation.

In all the above systems, the illusion is not perfect – users do not feel that the virtual objects are real. Despite this, users clearly are getting some indication of depth. What exactly are they perceiving?

For some of the applications mentioned above, such as advertising and weather systems, it is not important to portray an accurate impression of depth. For others, such as medical visualisation, it is imperative. A doctor needs to be sure that what he or she examines in the virtual world is a faithful representation of the patient's body. In the case of tele-surgery, the doctor must be able to accurately judge the distance between the tip of a scalpel and a vital organ.

We have been developing a FTVR system similar to Wedge [8] with the aim of creating a virtual sculpting tool. If the artist or designer does not perceive the virtual object as a

natural part of the 3D world then the FTVR interface may hinder their work rather than assist it. Thus, we are interested in evaluating the the success of the FTVR illusion.

In this paper, we review key issues regarding differences between the virtual world and the real world. We discuss the results of two previously published perceptual studies and a third of our own. We show that we still have an incomplete understanding of perception in FTVR and that some important questions may be impossible to answer. Although FTVR clearly is useful, it should be used with caution if the task requires accurate and immediate 3D perception.

2 Fundamental differences between FTVR and the Real World

There are a number of differences between the real world and FTVR. These are listed below:

- Fixed resolution - Fish tank virtual reality systems represent the virtual world using only a finite set of pixels. The causes a noticeable reduction in visual richness, as well as some unusual perceptual effects. In the real world, we move closer to an object to see it in more detail. When we do this the object takes up a larger amount of our field of view, and the size of the object's image on the retina of the eye increases. This action meets with strange results in FTVR. If we move closer to a virtual object that is located behind the screen, then the size of the retinal image increases in the same way as the real world, but the size of the object on the screen actually decreases. This results in the object appearing *less* detailed as it is drawn with fewer pixels. Virtual objects that are located in front of the screen behave in the opposite manner. This matches our natural expectations better, but it can still look strange. When approaching a small object located in front of the screen it will appear to gain too much resolution [4].

- Accommodation - The accommodation depth cue accompanies a change in the focal length of the lens in the eye [16]. When viewing a distant object, the muscles controlling the lens are relaxed and the lens is thin. Whereas when a close object is viewed, the eye accommodates by thickening the lens to focus the image. The state of the muscles that control the lens give a rough indication of depth. It is impossible to simulate the accommodation cue in FTVR because, in reality, the observed objects lie on the surface of the screen. If the virtual objects lie away from the screen, then the accommodation cue will conflict with the perspective depth cue manufactured by the system.

- Depth of field - The object you are looking at is normally in focus, while foreground and background objects appear out of focus [1]. In binocular FTVR systems it may be possible to measure, or estimate, the convergence of the gaze directions of the eyes and incorporate the depth of field cue [3], but it is currently not practical in real time systems.

- Lag - Because eye detection and rendering cannot happen instantly, there is a certain amount of lag inherent in all FTVR systems. This can be reduced by using predictive eye-tracking [3].

- Clipping - When the user moves too far to one side of the display, *clipping* of virtual objects occurs. When this happens it appears as though the edge of the screen

obscures the object and the illusion that an object is in front of the screen is destroyed. A partial remedy is to increase the user's field of view, by using more or larger screens.

- Edge effects - The edges of the display screen provide the user with a strong indication of depth. Virtual objects that lie in front of the screen and appear close to the edges of the screen are subject to conflicting depth cues. This is because the edges provide a stronger depth cue than the simulated perspective. In addition, binocular systems suffer from the problem that parts of objects near the edges will be visible to one eye but not the other. Many users find this objectionable and often perceive the object to be located behind the display screen [12]. This confusion can be reduced by placing a virtual border in front of the virtual 3D scene [13, 12].

- Interaction with the virtual world - Being physically present within a FTVR system can cause depth cue conflicts by occlusion. For example, reaching out to touch a virtual butterfly will most likely occlude part of the image and destroy the 3D effect. Even if the user's hand does not overlay any part of the image, it will provide conflicting depth cues in a similar way to the edge effect.

- Distortion in binocular FTVR - Users of binocular systems have trouble fusing the stereo image pairs. The common solution to this problem is to intentionally underestimate the eye-separation of the user [20]. However this creates a new problem with head-tracked FTVR displays: the stereoscopic image will appear to shift and warp with the head movement. While it is impossible to completely remove this distortion, a partial correction is available [20].

All of these differences provide the user with clues that the virtual world is not real and they make it difficult to compare perception in the virtual world with perception in the real world.

3 Practical Difficulties in FTVR Systems

In addition to the theoretical limitations listed in the previous section, there are a number of practical difficulties in setting up a FTVR system. A poor implementation will disrupt the illusion of reality even further.

A large potential source of error is the misregistration of the virtual world with the real world. Registration information is captured in the perspective viewing matrix, defined by a frustum whose tip is located at the optical centre of the eye and whose base is defined by the four corners of the display surface [3, 1]. Stationary virtual objects in an incorrectly registered system will appear to move as the viewpoint changes.

In order to define the viewing matrix accurately, the position of the optical centre of the eye must be known. Since the eye's optical centre is not in the same place as its rotational centre [3], gaze direction must be measured in addition to pupil location. Most FTVR tracking systems do not account for gaze direction [19, 1]. Some do not even track the pupils: They track other markers on the head and estimate the pupil location as a fixed offset from the markers [5].

Other practical difficulties include:

- Brightness discontinuities - Using more than one display screen can dramatically improve the viewing area, particularly if the screens surround the user or surround the virtual scene. Since most displays do not radiate light equally in all directions, brightness discontinuities are likely to appear at the boundaries between the screens. This discontinuity is distracting to the user and may even affect the 3D illusion [5]. While it is not possible to completely remove the discontinuity in software, it is possible to improve the situation [17].
- Pixel appearance - When viewing the pixels of some displays they look like small square patches of light, surrounded by a thin black border. This gives the image a "grainy" texture. We have noticed that this effect is especially obvious with LCD projector displays. This could be corrected by defocusing the projector or through improved display technology.
- Refractive and curvature distortions of glass - CRT monitors are viewed through a layer of glass. Glass or transparent plastic may also be used in rigid backprojection systems. In both cases refraction will redirect light as it travels towards the user, so corrections must be made for the refractive index of the material. Curvature of the glass in CRT screens introduces additional distortion, while projector systems are relatively distortion free [1].

4 Designing Perception Experiments

There are several factors that need to be considered when designing perception experiments. These include the following:

- Registration - One cannot hope to perform perception experiments unless the user is being presented with the correct image. This can only be achieved if the virtual world is registered correctly with the real world.
- Switching - Some of the experiments described in this paper require the subject to frequently change focus between the virtual world and the real world. The effect of this on performance is uncertain.
- Deliberation - In order to test perception, a task is given to a subject. Performance may be affected if the subject takes time in making a judgement rather than using an instinctive approach.
- Cheating - In trying to solve the task given, some subjects seek and find ways to perform the task by using their knowledge of geometry. For example, in our system the virtual object could be located by viewing it from several planned angles. In this case, the subject may produce a good result even without the illusion of depth.
- Individuality - In order to test perception, we require subjects to perform a perception task. It is not possible to separate the subject's perception from the subject's ability to perform the task. Furthermore, individual subjects may perform differently, for reasons such as different motor skills.

5 Perception Experiments

Perception within FTVR systems has not been studied extensively [5]. This lack of a comprehensive amount of material seems to extend to the wider field of VR systems, although some work been done [18, 6].

We evaluate two previously published studies of FTVR with respect to the factors mentioned in Section 4, and we discuss a study of our own.

5.1 Djajadiningrat et al. (1998)

The designers of Cubby [5] performed a detailed experiment in perception. A subject was presented with a pyramid and a cube, either real or virtual, and was asked to estimate the distance between the tip of the pyramid and the centre of a marked face of the cube. They adjusted a blank measuring instrument, similar to callipers, to match this distance. This equipment is shown in Figure 2.

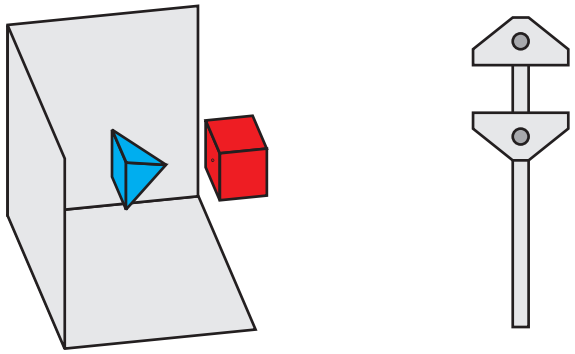

Figure 2: (a) The perception experiment performed in Cubby. (b) The measuring instrument.

The experiment was performed in Cubby monocularly, and the results compared with both the subject's monocular and binocular perception in the real world. It was found that performance in Cubby could not match that in a binocularly viewed real scene, but it compared well to a monocularly viewed real scene. More specifically, the average error in Cubby was similar to the average error in the real scene, but the spread of the error was larger in Cubby.

An advantage of the experiment done in Cubby is that the subject's ability to perform the task did not greatly affect the result. Less able subjects could take more time until they were satisfied with the adjustment of the measuring instrument.

The disadvantage of this task is that it requires the subject to switch attention between the virtual and the real worlds. There is also a significant amount of deliberation required to do the task.

An infrared tracker was used to locate a marker on the face rather than the eyes. The eye position could change with head rotation so the system was subject to registration errors.

5.2 Yoshida et al. (1998)

Some experiments were performed to test a subject's perception of depth in a fixed-position, binocular VR system [22]. While not strictly a FTVR system, because the position of the viewer was fixed, the experiments are relevant to our discussion. The fixed viewpoint makes it easier to register the system accurately.

For this experiment a real or virtual cone was placed at various distances from the screen. The subject then reached into the world and indicated, with a tracked pointer, the location of the tip of the cone. The difference was calculated and the results compared. They showed that there was a difference between the task in the real world and the task in the virtual world. Furthermore, they indicated that there may be some difference in performance between subjects. They used these results to devise a method for correcting for the difference between the virtual world and the real world.

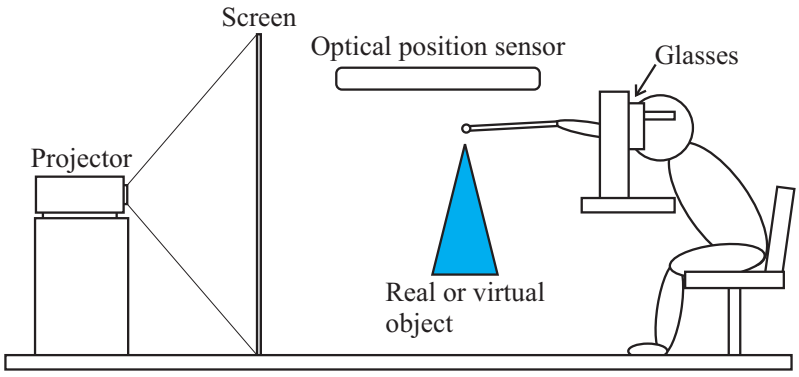

Figure 3: Measuring depth in the virtual world.

The correction procedure used a false eye-separation distance. As stated earlier, this creates a distortion in the image. However, it would not have been noticeable in their experiments because of the lack of head-tracking.

This experiment involves placing a real object into the virtual world, some attention switching between the virtual and the real worlds, and some deliberation.

5.3 Treadgold et al. (2000)

We set out to design an experiment that would not require the user to switch his or her visual attention between the virtual and real worlds.

Our task is an abstraction of the process of catching a falling object. People can usually do this without having to see their hands. A cube, real or virtual, was suspended above two opaque surfaces. The user places his or her hand underneath the top opaque surface, and must indicate on the bottom surface where the cube would land if it were to fall straight down.

5.3.1 Experiment 1 - Monocular and Binocular Perception in the Real World

We first wanted to try the experiment in the real world to see how accurately people could perform the task. At the same time, we compared monocular and binocular perception.

For this experiment, we constructed the equipment shown in Figure 4. The top platform was made of a transparent plastic sheet, marked with a regularly spaced 2cm grid. The grid was 22 squares wide and nineteen squares deep. The middle platform was made of a piece of thick card that prevented the subject from seeing where he or she placed the marker. The bottom platform was covered in a lined grid that matched the grid on the transparent layer.

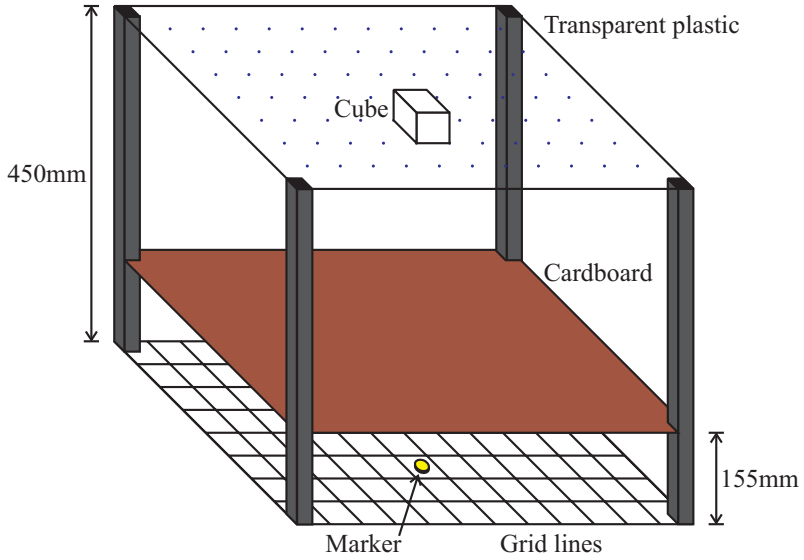

Figure 4: The equipment used to test monocular and binocular perception.

For each trial of the experiment a cube was placed on the plastic layer. The position of the cube was determined by a random number generator and the sequence was different for each subject. The subject was then asked to place a marker underneath the position of the cube. The position of the marker was recorded and the trial repeated.

Each subject completed two sets of trials; one done monocularly, and the other done binocularly. Each set consisted of 35 trials, which included an initial five practice trials. A break of at least a day was given between sets in order to minimise any learning effects. In total, ten subjects participated in this experiment.

We calculated the *error* in each trial by computing the difference, in the horizontal plane, between the position of the cube and the position of the marker. We averaged the error for each subject and compared the monocular task with the binocular task. These results can be seen in Figure 5. The graph shows the average error for each subject in the monocular task and the binocular task, as well as the standard deviation of the errors.

The graph shows that there is a reasonable difference in the size of the error between subjects. To analyse the difference we first consider the monocular task, and then the binocular task.

To determine whether the difference was significant, we performed an analysis of variance (ANOVA) test on the subject group with the null hypothesis that the mean error for all subjects, while doing the monocular task, is the same. We found that there is very good reason to believe that there is a difference between the means of at least two of the subjects ($F = 25.129; df = 9; p < 0.000$).

To determine the difference between subjects we followed with the Tukey test, which compares each subject with every other subject. The results of this are shown in Table 1. The table shows that the subjects can be divided up into four subgroups. Within each subgroup we are confident that, within the 95% confidence level, there is no significant difference between the means of the subjects. For example, looking at Column 1, there is no significant difference between Subjects 6, 1, 8, 9, 10, and 7, but because of the overlap with Column 2, there is also no significant difference between Subject 7 and Subject 2.

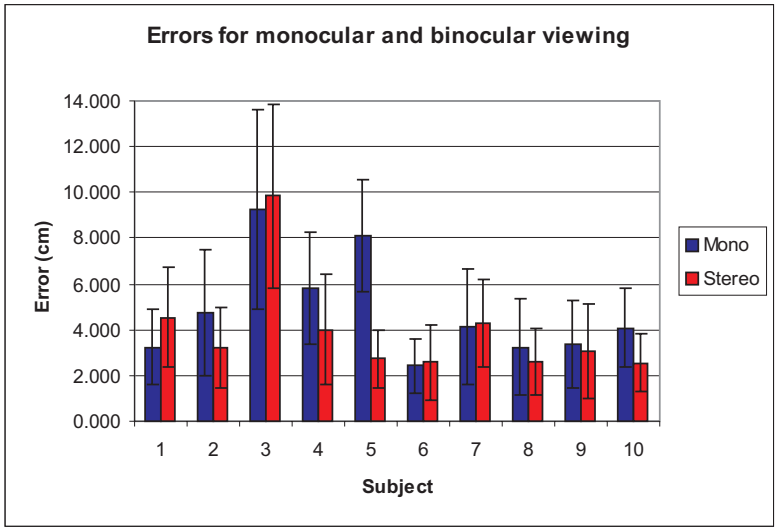

Figure 5: The errors in the monocular and the binocular viewing conditions.

There is however, a significant difference between Subject 2 and Subject 6 because they do not appear in the same column. Because there is no overlap between Column 4 and the other columns, the subjects in this column, Subjects 5 and 3, are significantly different from all the other subjects.

Table 1: The Tukey test for the monocular task.

	Subset for 95% confidence level			
subject	1	2	3	4
6	2.413			
1	3.236	3.236		
8	3.247	3.247		
9	3.381	3.381		
10	4.075	4.075	4.075	
7	4.132	4.132	4.132	
2		4.735	4.735	
4			5.816	
5				8.100
3				9.228

We performed the same tests for the binocular task. In a similar way, we found that there was strong evidence to suggest that there is a difference between the means of at least two of the subjects ($F = 31.826; df = 9; p < 0.000$). The Tukey test also revealed a table similar to the monocular task. It is not shown here because we have already established that subjects perform differently from each other.

We now know that the performance is dependent on the subject. We therefore compared the monocular task with the binocular task for each individual subject, rather than comparing the group as a whole.

To do this we performed a two-tailed, unpaired t-test on each subject. This test assumes that the samples are uncorrelated; that each trial of the task does not affect performance in subsequent trials. We believed that the correlation within the samples would be low

Table 2: Comparing monocular viewing with binocular viewing.

Two-tailed, unpaired t-test at the 95% confidence level; critical value = 1.96										
Subject	1	2	3	4	5	6	7	8	9	10
t-value	2.62*	2.51*	0.58	2.91*	10.6*	0.49	0.32	1.44	0.62	3.94*
Error	↑	↓	-	↓	↓	-	-	-	-	↓
*above the critical value										

because there was no feedback during the experiments. The subject simply placed the marker underneath the cube. They could not see where it was placed, and they were not told how far off they were. These conditions make it difficult to improve at the task over the course of the experiment. Indeed, the results of an autocorrelation test, which checked over the first ten lags, showed that there was no correlation within the samples.

The null hypothesis for the t-test is that the mean error in the monocular is equal to the mean error in the binocular task. The result of the t-tests are shown in Table 2. It shows that five out of the eleven subjects yielded a t-value below the critical value. For these subjects the null hypothesis is valid, and there is no difference between the mean error in the monocular task and the mean error in the binocular task.

For five of the ten subjects we must reject the null hypothesis and acknowledge that there is a difference between the mean errors. The fourth row in Table 2 indicates whether the use of binocular vision reduced or increased the error. The performance of four of these five subjects increased with the use of binocular vision, while the performance of the remaining one decreased.

We found that performance of the task was dependent on the individual.

Our results showed that performance of half of the subjects was the same in the monocular task as the binocular task. Of the other half, the performance of four of the subjects increased with the use of binocular vision, while the performance of the of the other one decreased.

There is some evidence to suggest that performance may be hindered by doing the task monocularly. This result is contrary to that of Djajadiningrat et al. (1998), who found that there was no significant difference between the performance in their monocular task compared with their binocular task.

5.3.2 Experiment 2 - Monocular Perception in the Real and Virtual Worlds

A similar setup to the first experiment was used, except that it was made triangular to conform with the shape of our FTVR system. The spacing of the grid on the bottom platform was 1cm, and the spacing of the grid on the top platform was the larger size of 3cm. The triangular grid on the top platform forms an isosceles triangle with an equal-side-length of thirteen squares. This setup is shown in Figure 6.

Each subject completed two sets of 30 trials, including five practice trials. One set of trials were made in the real world and the other were made in the virtual world. Half of the subjects performed the real task first, and the other half performed the virtual task first. This was done to even out any effects that could be caused by doing one or other of the sets first. In total, fourteen subjects performed this experiment.

We matched the virtual world as closely as possible to the real world by shading the cube and simulating the grid marked on the transparent plastic layer with a layer of virtual blue dots. These environments are shown in Figure 7.

After the experiment we conducted an informal interview with each of the subjects.

Out of the fourteen subjects we removed the results of two of them because they used a technique in the virtual world that we considered "cheating". They would view the cube along an axis that was perpendicular to one side of the Wedge and which passed through the centre of the cube, and then do the same thing on the other side. We considered this cheating because they were using their knowledge of geometry, rather than trusting their intuitive perception of depth.

The graph in Figure 8 shows the errors in the real and the virtual worlds. After checking the correlation, we again performed an unpaired t-test for each subject in the experiment. The results of this are shown in Table 3. The t-values indicate that there was no significant difference in the mean error for three out of the twelve subjects. There was a significant difference for the other nine subjects. The fourth row of this table shows whether the mean error was larger or smaller in the virtual world compared with the real world. Six of the subjects had a higher mean error in the virtual world, and three of the subjects had a lower mean error in the virtual world.

The results of the informal interview are summarised below:

- All subjects reported that the cube appeared to be in front of the screen.

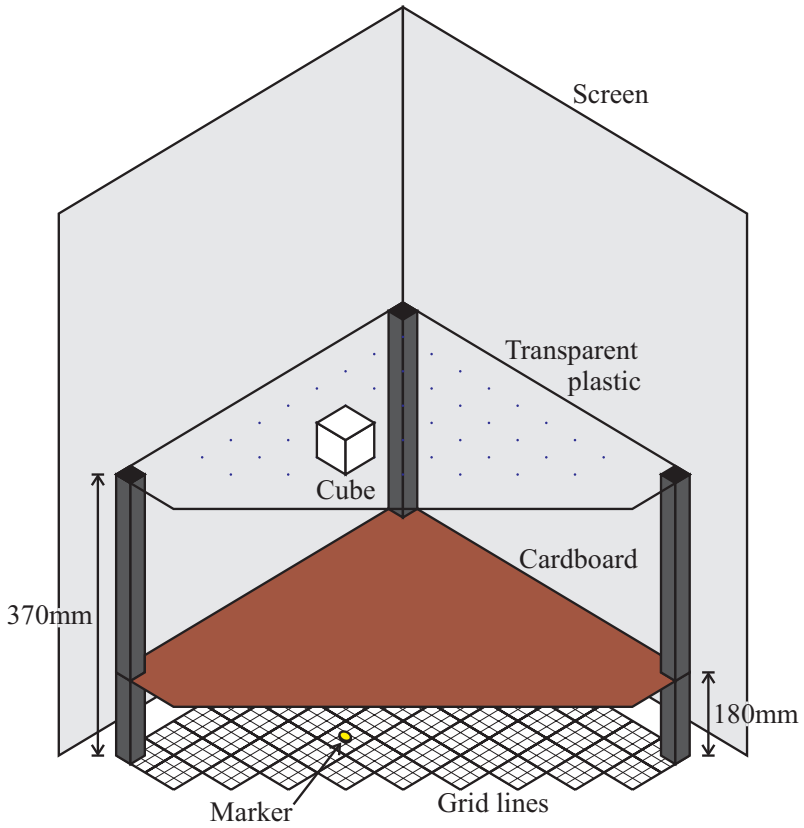

Figure 6: The equipment used to test real and virtual perception.

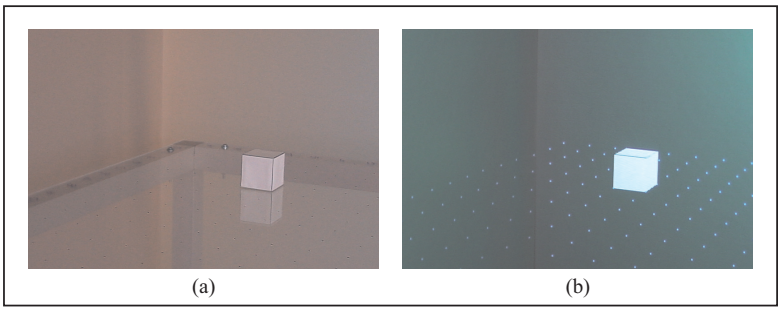

Figure 7: (a) The real environment. (b) The virtual environment.

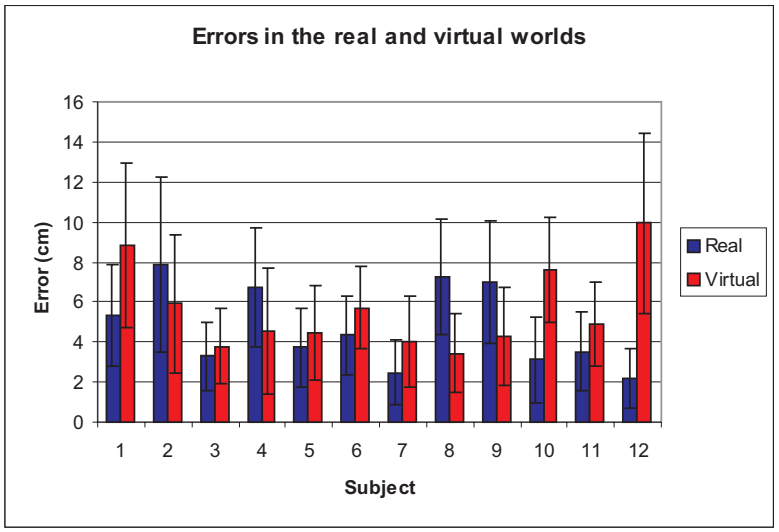

Figure 8: The errors in the real and the virtual worlds.

Table 3: Comparing the real world with the virtual world.

Two-tailed, unpaired t-test at the 95% confidence level; critical value = 1.96												
Subject	1	2	3	4	5	6	7	8	9	10	11	12
t-value	6.30*	1.70	0.94	2.38*	1.20	2.34*	2.76*	5.04*	3.23*	6.25*	2.33*	7.78*
Error	↑	-	-	↓	-	↑	↑	↓	↓	↑	↑	↑
*above the critical value												

- A few subjects mentioned that the cube seemed to move around as they moved their head.
- One subject commented that the cube appeared to be always located just in front of the screen, and that it appeared to get bigger rather than appearing closer to the eye.

Matching the virtual world with the real world turned out to be more difficult than expected and the difference was very noticeable.

For the first part of this experiment, we found that nine of the twelve subjects responded differently to the virtual task compared with the real task. The performance of six of these subjects was reduced in the virtual world compared to the real world, while the performance of the other three increased. The results indicate that different subjects have

an indeterminable response to the virtual task compared with the real tasks.

From the informal questionnaire we can be sure that the 3D illusion was perceived. The subjects who commented that the cube appeared to move were probably observing the inaccuracies in the registration. The subject who commented that the cube appeared to get bigger, rather than closer, pointed out a problem that we had also noticed. We suspect that because size gives a strong indication of depth, it is taking precedence over the motion parallax.

5.4 Summary

The results of the three perception experiments are summarised in Table 4.

Table 4: Summary of the three perception studies.

	Djajadiningrat et al. (1998)	Yoshida et al. (1998)	Treadgold et al. (2000)
FTVR system	3 screens, monocular, head-tracked	1 screen, binocular, fixed head position	2 screens, monocular, head-tracked
Registration	3 DOF tracking	fixed head position	5 DOF tracking
Switching	yes	yes	no
Deliberation	yes	yes	yes
Individuality	no[1]	yes	yes
Real world comparisons	monocular & binocular	binocular only	monocular and binocular
Results	monocular FTVR similar to monocular real, monocular FTVR not as good as binocular real	binocular VR not as good as binocular real	monocular real not as good as binocular real, indeterminable results when comparing FTVR with real

[1]Not mentioned in results. Good reason to believe that individuality is minimal.

6 Conclusions and Future Work

There is a lot left to be discovered about perception in FTVR. A complete understanding will be difficult to achieve since we cannot measure perception directly; only through performance on tasks related to perception. There are three key difficulties:

- Measuring immediate perception - We want to know that our subjects are getting an immediate and natural perception of virtual space. However, some subjects use high-level reasoning and knowledge of geometry to perform the task. Furthermore, it is difficult to separate perception from reasoning. Demanding a quick response may help.
- Varying skill levels - If subjects vary in their skill level at performing the task, then generalisation of the results for the population as a whole is difficult. Furthermore, individuals should be able to perform the task with consistent accuracy. One approach towards reducing this problem would be to use subjects who are experts in understanding spatial relationships in the real world, such as artists and designers. This may be a particularly useful approach for us since artists and designers will be our end users.

- Providing an accurate FTVR system - As discussed in Section 4, some problems in FTVR are fundamental and others are dependent on the implementation. If we want to make conclusions about FTVR systems in general, it is essential that the implementation is as good as possible. The biggest implementation issue is that of registration. Most existing FTVR systems, including our own, simply do not have the tracking technology to locate the optical centre of the eye accurately.

Our user study was an advance on earlier studies in that we were able to measure performance without requiring subjects to switch their visual attention between the real and virtual world. However we were not able to satisfy the criteria described above.

FTVR has proven its value for games and for scientific visualisation. User studies have shown that it is particularly useful for determining the depth ordering of objects in the virtual world [19]. However, we still do not know how accurate users' spatial perception of the virtual world is. Until this question is settled, we cannot conclude that FTVR is well suited for crucial applications such as 3D model design and medical simulation.

7 Acknowledgements

We would like to thank Susan Galvin from the University of Otago Department of Psychology for her advice on perception experiments.

This project was supported by a grant from the Marsden Fund.

References

[1] C. Cruz-Neira, D. J. Sandin and T. A. DeFanti: Surround-Screen Projection-based Virtual Reality: The Design and Implementation of the CAVE in: *Computer Graphics (SIGGRAPH '93 Proceedings)*, ed. J. T. Kajiya vol. 27 (Aug. 1993) pp. 135–142.

[2] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon and J. C. Hart: The CAVE: audio visual experience automatic virtual environment **35** (1992)(6) 64–72.

[3] M. Deering: High resolution virtual reality *Computer Graphics* **26** (1992)(2) 195–202.

[4] M. F. Deering: Making virtual reality more real: experience with the virtual portal in: *Proceedings of Graphics Interface '93* Canadian Information Processing Society, Toronto, Ontario, Canada (May 1993) pp. 219–226.

[5] J. Djajadiningrat: *Cubby: What You See Is Where You Act* Ph.D. thesis Technical University of Delft, The Netherlands (1998).

[6] G. W. Fitzmaurice, S. Zhai and M. H. Chignell: Virtual reality for palmtop computers *ACM Transactions on Information Systems* **11** (1993)(3) 197–218.

[7] R. H. Fowler, W. A. L. Fowler, A. Kumar and J. L. Williams: Head-coupled Stereo Display for Visualization in a Document Retrieval System using Associative Networks (Jul. 1995) Online. Accessed 1 December 2000.
`http://www.cs.panam.edu/info_vis/sigir95.html`

[8] H. J. Gardner, R. W. Boswell and D. Whitehouse: The WEDGE Emmersive [*sic*] Projection Theatre in: *The Simulation Technology and Training (SimTecT99) Conference* (Mar. 1999) .

[9] W. Kruger, C.-A. Bohn, B. Frohlich, H. Schuth, W. Strauss and G. Wesche: The responsive workbench: A virtual work environment *Computer* **28** (1995)(7) 42–48 ISSN 0018-9162.

[10] J. Liang and M. Green: JDCAD: A Highly Interactive 3D Modeling System **18** (1994)(4) 499–506.

[11] C. Maggioni and B. Kämmerer: GestureComputer - History, Design and Applications in: *Computer vision for human-machine interaction*, eds. R. Cipolla and A. Pentland Cambridge; New York : Cambridge University Press (1998) pp. 23–51.

[12] J. D. Mulder and R. van Liere: Enhancing Fish Tank VR in: *IEEE Virtual Reality 2000 Conference Proceedings* (March 2000) .

[13] S. Ohtsuka, Y. Ishigure, Y. Kanatsugu, T. Yoshida and S. Usui: Virtual Window: A Technique for Correcting Depth-Perception Distortion in Stereoscopic Displays *The Society for Information Display Symposium Digest* **27** (1996)(893).

[14] C. Regan: An investigation into nausea and other side-effects of head-coupled immersive virtual reality *Virtual Reality: Research, Develpment and Application* **1** (1995)(1) 17–31.

[15] J. Rekimoto: A Vision-Based Head Tracker for Fish Tank Virtual Reality – VR without Head Gear – in: *IEEE Virtual Reality Annual International Simposium (VRAIS)* (1995) pp. 94–100.

[16] R. Sekuler and R. Blake: *Perception* McGraw-Hill, Inc 3rd edn. (1994).

[17] M. Treadgold, K. Novins and G. Wyvill: Correcting Brightness Discontinuities in Fish Tank Virtual Reality Systems in: *Image and Vision Computing New Zealand 2000* (Nov. 2000) pp. 84–89.

[18] L. R. Wanger, J. A. Ferwerda and D. P. Greenberg: Perceiving spatial relationships in computer-generated images *IEEE Computer Graphics and Applications* **12** (1992)(3) 44–58.

[19] C. Ware, K. W. Arthur and K. S. Booth: Fish Tank Virtual Reality in: *Proc. ACM Conf. Human Factors in Computing Systems, INTERCHI (INTERACT & CHI)*, eds. S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel and T. White ISBN 0897915755 (24–29 Apr. 1993) pp. 37–42.

[20] Z. Wartell, L. F. Hodges and W. Ribarsky: Balancing fusion, image depth and distortion in stereoscopic head-tracked displays *Computer Graphics* **33** (1999)(Annual Conference Series) 351–358 ISSN 0097-8930.

[21] K. Waters, J. Rehg, M. Loughlin, S. Kang and D. Terzopoulos: Visual Sensing of Humans for Active Public Interfaces in: *Computer vision for human-machine interaction*, eds. R. Cipolla and A. Pentland Cambridge; New York : Cambridge University Press (1998) pp. 83–96.

[22] S. Yoshida, T. Hoshino, T. Ozeki, S. Miyazaki, J. Hasegawa, T. Yasuda and T. Fukumura: An Improvement Method of Depth Perception in Stereoscopic Display *International Archives of Photogrammetry and Remote Sensing* **32** (1998) 386–391.