Changes in Navigational Behaviour Produced by a Wide Field of View and a High Fidelity Visual Scene

S. Lessels and R. A. Ruddle

School of Computing, University of Leeds. UK

Abstract

The difficulties people frequently have navigating in virtual environments (VEs) are well known. Usually these difficulties are quantified in terms of performance (e.g., time taken or number of errors made in following a path), with these data used to compare navigation in VEs to equivalent real-world settings. However, an important cause of any performance differences is changes in people's navigational behaviour. This paper reports a study that investigated the effect of visual scene fidelity and field of view (FOV) on participants' behaviour in a navigational search task, to help identify the thresholds of fidelity that are required for efficient VE navigation. With a wide FOV (144 degrees), participants spent significantly larger proportion of their time travelling through the VE, whereas participants who used a normal FOV (48 degrees) spent significantly longer standing in one place planning where to travel. Also, participants who used a wide FOV and a high fidelity scene came significantly closer to conducting the search "perfectly" (visiting each place once). In an earlier real-world study, participants completed 93% of their searches perfectly and planned where to travel while they moved. Thus, navigating a high fidelity VE with a wide FOV increased the similarity between VE and real-world navigational behaviour, which has important implications for both VE design and understanding human navigation.

Detailed analysis of the errors that participants made during their non-perfect searches highlighted a dramatic difference between the two FOVs. With a narrow FOV participants often travelled right past a target without it appearing on the display, whereas with the wide FOV targets that were displayed towards the sides of participants overall FOV were often not searched, indicating a problem with the demands made by such a wide FOV display on human visual attention.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques-Interaction techniques. I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism-Virtual Reality. H.5.2 [Information Interfaces and Presentation]: User Interfaces-Input devices and strategies

1. Introduction

It is well known that people frequently have difficulty navigating in virtual environments (VEs), and take substantially longer to develop spatial knowledge of such environments than equivalent real-world settings [WBKP96] [RMH99]. The root cause of this difficulty is assumed to lie in a reduction in the fidelity of VEs compared with the real world, with this fidelity classified along independent dimensions such as the content of the visual scene, the field of view (FOV) used to observe that scene, and the mechanism used for movement [WHK98].

Navigational difficulties are found in a variety of different types of VE, from seascapes covering hundreds of square kilometres, to sections of virtual countryside, virtual buildings and even single virtual rooms [DS96] [RMH99] [RJ01] [SM00] The magnitude of these difficulties, and any comparison with the real world, is usually measured in terms of differences in users' navigational performance (e.g., distance travelled, time taken or errors made). However, it is also important to compare users' navigational behaviour, because subtle changes in that behaviour can lead to a drastic deterioration in navigational performance. In other words, it is behaviour that often dictates performance. If a VE could be designed so that users behaved in a similar manner to the

real world then there should be a large improvement in navigational performance, bringing substantial benefits to a wide range of VE applications.

This paper describes an experimental study that investigated participants' navigational behaviour when they navigated a virtual room, searching for targets that had been placed in locations that were explicitly marked. The task was known to be straightforward to complete in the real world but non-trivial in a VE [LR03] The purpose of the experiment was to investigate any changes that occurred in navigational behaviour when the virtual scene was changed from low to high fidelity and the FOV increased from 48 (normal for a VE) to 144 degrees (wide), and compare both behaviour and performance to the earlier real-world study. First, however, we outline previous research into the effects of visual fidelity and FOV on spatial tasks.

1.1. Visual fidelity and FOV

Waller, Hunt, & Knapp [WHK98] included visual fidelity within their concept of environmental fidelity, which is how closely a VE resembles its corresponding real world scene. The content of the visual scene is one of the primary sources of information that people use to determine their position and orientation within a VE. It follows, that changing the amount of visual information available will affect users behaviour and that will, in turn, affect their ability to perform spatial tasks in a VE. Information regarding position and orientation in the VE can be made available via texture, including landmark objects embedded within textures. The amount of that information that is visible at any single moment in time is dictated by the FOV.

To navigate successfully participants must perform some degree of distance estimation either from their movement (path integration) or by updating their position and orientation from the visual scene, or a combination of both. The repetition (tiling) of certain textures, such as a brick pattern, has been shown to significantly improve participants' accuracy when estimating distances [SKD*99]. However, distance estimations were not improved with the use of grass, carpet, or abstract patterned textures [WK98] [SKD*99]. Textures can also enhance optic flow, which has been shown to aid path integration, prevent disorientation, and helps participants locate remembered target positions more accurately. [KKSB00] [KWDT02]. Texture maps are a cost effective method for improving the visual realism of VEs. These types of textures often either are, or contain, distinctive "objects" that act as landmarks and eliminate the need for them to be modelled in 3D geometry. The saliency of landmark cues has also been shown to influence navigational performance. In this respect, participants' route-finding accuracy has been shown to increase when everyday (i.e. familiar) objects were used as local landmarks, but not when abstract (i.e. unfamiliar) patterns were used instead [RPJ97].

Studies that have investigated the effect of FOV in VEs have often suggested that a restricted FOV contributes toward poor spatial performance. For example, when participants judged egocentric distances while wearing a HMD, they underestimated the distances while using a wide FOV (140 x 90 degrees) but overestimated the distances while using a narrower FOV (60 x 38.5 degrees; [KW96].

Present research concentrates primarily on performance measures very little research has attempted to identify and measure users underlying behaviour. However, a recent study by Czerwinski, Tan and Robertson (2002) [CTR02] required participants to find target objects and move them to predefined locations as quickly as possible. With a novel navigation technique (the faster participants moved forward the higher and steeper they viewed the VE) males performed faster then females with a narrow FOV, but with a wide FOV females' speed increased to equal that of males. Data analysis was not only performed on performance results (speed) but also on how participants achieved those results. The behavioural metric showed that males travelled higher than females with the narrow FOV, while both genders travelled higher with the wide FOV. Also in the wide FOV there were differences in the use of strategy between the genders, females travelled less distance to achieve a faster time, while males travelled further, but quicker, to achieve a faster time. These behavioural metrics show the different strategies adopted and illuminate how performance results occurred.

There are also VE studies that have shown no influence of FOV for spatial tasks. In one, participants performed a series of triangle completion tasks using a large projection screen, and accuracy was influenced by path layout but not by FOV [PMW97]. Another investigated participants' ability to estimate ego-rotations [SPRvdHB02]. There was no difference between two FOVs presented on a projection screen (86 x 64 degrees, and 40 x 30 degrees), but with both of these participants were significantly more accurate than when the judgments were performed using a 40 x 30 degree HMD.

2. Experiment

The experiment investigated participants' performance and navigational behaviour when they searched a room-sized VE that contained a visual scene constructed using one of two fidelities, and was displayed using one of two FOVs. The visual scenes were a photorealistic model of a sports hall (high fidelity) and a model with identical geometry but which was rendered using a tiled brick texture (low fidelity). The FOVs were normal (48 degrees; displayed on a single computer monitor) and wide (144 degrees; three monitors arranged in an arc that subtended the same angle). A 2 by 2 between participants design was used, with each participant searching the VE while using one combination of scene fidelity and FOV.

2.1. Method

2.1.1. Participants

Forty two participants (22 females and 20 males) took part in the experiment and their ages ranged from 21 to 36 (M = 24.35, SD = 3.07). Two female participants suffered from nausea and withdrew from the experiment. Presented here are the results of the 40 participants who successfully completed the experiment. All the participants volunteered for the experiment and were paid an honorarium for their participation.

2.1.2. Materials

Participants travelled around the VE using a keyboard and mouse to interact. Participants could look up and down, and travel forward in the direction they were looking (view-direction movement). When a participant held down the keyboard 'up' cursor they moved forward at a speed of 1 m/s and stopped when the key was released. The mouse was used to control the participant's direction of view. Looking up and down was achieved using zero order control; by moving the cursor up the screen the viewing pitch increased by up to +90 degrees, and the viewing pitch decreased by a corresponding amount if the cursor was moved down the screen. Looking left and right was accomplished using first order control with the rate of turning increasing proportionally with the cursor's distance away from the vertical centreline of the screen. The maximum rate of turning was 135 degrees/second.

The VE contained 33 cylinders, all measuring 0.5m in diameter and 1.35m high. Thirty-two were arranged in eight identical groups of four, while the 33rd cylinder was positioned in the centre (Figure 1).
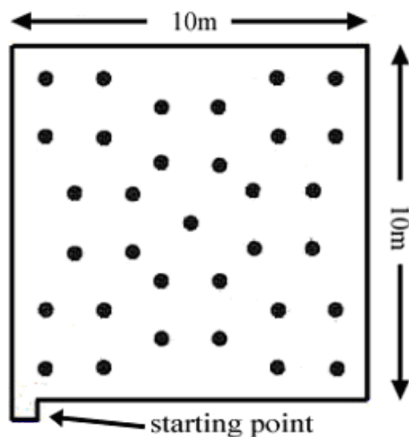

Figure 1: *Plan of the VE.*

On top of 16 of the cylinders was a white box with a

blue lid that indicated a possible target location. Eight of the boxes contained a target object (a red square) while the remaining eight were empty (decoy boxes). To ensure that targets and decoys were distributed around the environment one box in each group of four cylinders was randomly chosen in each trial to be the target and another a decoy.

Participants could raise and lower the lids of the target and decoy boxes by pressing the left mouse button. Pressing it once raised the lid, and pressing it again lowered the lid. If there was a target present inside the box the participant pressed the right mouse button to select it, the target then turned from red to blue indicating that it had been found and the VE software then automatically lowered the lid. The software prevented participants from moving away from any box until its lid was lowered. A lid could only be raised, and a target selected, if the participant was within 0.6 m of the centre of the box (i.e., they were adjacent to it) and it was within the participant's FOV.

The high-fidelity scene condition used images of the sports hall surfaces as textures, these were captured using a digital camera and 'stitched' together to create seven separate textures: each of the 4 wall textures were 1024 x 512 pixels (1.5 Mb each in RGB format); the ceiling texture was 512 x 128 pixels (0.4 Mb). One texture was used for each of the 25 lights suspended from the sports hall ceiling (0.1 Mb each). The seventh texture was used for the floor and consisted of a repeated base texture and was 2048 x 1024 pixels (6.1 Mb).

The low fidelity VE used the same floor and ceiling textures but a repeated brick texture that had no other features for each of the four walls. Thus in the high fidelity VE all four walls of the sports hall were clearly distinguishable from each other (Figure 2), but in the low fidelity (brick tiles) VE all four walls looked similar (Figure 3).


Figure 2: *The four walls of the virtual sports hall (high fidelity scene.*

In the normal-FOV condition participants viewed the VE via a monitor with a 475 mm x 300 mm viewable screen size. The resolution was 1280 x 1024 pixels and the horizontal FOV was 48 degrees. The wide-FOV conditions used the same monitor as the normal-FOV conditions but with two additional monitors of the same type and size placed
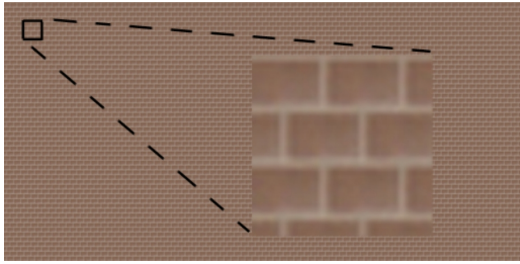
Figure 3: *Close-up of the brick texture that was repeated (tiled) across all four walls of the low fidelity VE.*

on its left and right. Each monitor displayed 1 of 3 viewing frustums at a resolution of 800 x 600 pixels. Taken together, the three monitors provided a continuous 144 degrees horizontal FOV of the VE, which was similar to their physical arrangement (Figure 4). The difference in the views seen by the participants with the two FOVs is shown in Figure 5.

The VE application was written in C++ and OpenGL Performer[TM] and ran on an SGI Onyx 3400 with a frame rate of 60 Hz. All participants viewed the monocular displayed VE from a distance of approximately 0.55m.

2.1.3. Procedure

Participants were randomly allocated to one of the four conditions, were run individually, and took approximately 45 minutes to complete the experiment. Each participant first practiced using the view-direction movement interface until they could fluently use the controls. The participant then performed two practice trials that allowed him or her to become familiar with the search task, and then completed four test trials. Each trial began at the starting point in the boundary recess (Figure 1) and participants searched until they had found and selected all eight targets. They were informed that the targets were always in the white boxes, but that their positions changed between trials. They were asked to complete the task as efficiently as possible by minimizing the distance travelled and to avoid revisiting possible target locations.

2.2. Results

Participants' performance in each trial was measured by recording the time they took to find the eight targets and by calculating how close participants came to conducting a "perfect search" (i.e. only checking each target and decoy once). Participants' behaviour was measured in three ways:

- by calculating the proportion of time spent performing different types of action
- classifying their search strategy.
- by assessing any errors they made that prevented them from conducting a "perfect search"
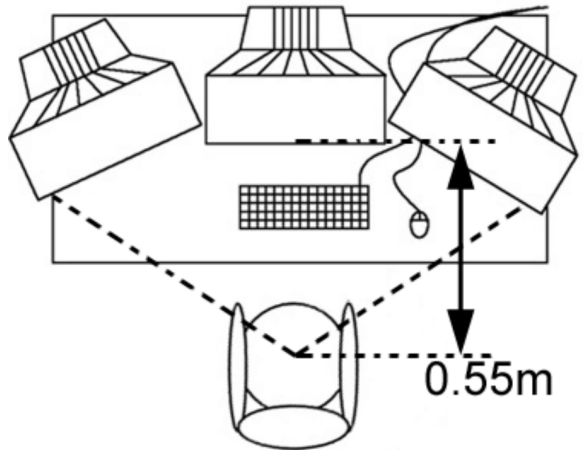
Figure 4: *Photograph of the 3-monitor setup (above), and plan (below).*

Statistical analyses of the data were performed using mixed design ANOVAs that treated the visual fidelity (brick tiles vs. sports hall) and FOV (normal-FOV vs. wide-FOV) as between participants factors, and the trial number as a repeated measure. Interactions are only reported if they were significant.

In terms of performance, participants took less time to find the targets with the sports hall than with the brick tiled scene but the difference was not significant (M = 168 seconds vs. 208 seconds, $F_{(1, 36)}$ = 2.21, $p > .05$). Participants also took less time to find the targets with the wide FOV than with the normal FOV, but again the difference was not significant. (M = 163 seconds vs. 212 seconds, $F_{(1, 36)}$ = 3.29, $p > .05$). However a significant learning effect occurred with participants taking less time to find the targets as the trials progressed ($F_{(1, 108)}$ = 3.70, $p < .05$). The mean time reduced from 218 seconds (trial 1) to 176 seconds (trial 4).

Figure 5: *Views from each frustum for the wide FOV condition, high fidelity (above), low fidelity (below). Only the middle frustum was visible in the normal FOV condition.*

When performing the task in the real world, participants completed 93% of the trials perfectly, visiting each target and decoy only once [LR03]. However, in the present VE experiment only 51% of the trials were completed perfectly, and it was noticeable that participants often only found six or seven of the targets before checking again at least one target or decoy that had already been visited. A repeated measures ANOVA of the number of targets found before any revisitation showed a significant interaction between FOV and scene fidelity ($F(1, 36) = 5.24$, $p < .05$), with participants finding most targets with the wide FOV/sports hall and least with the wide FOV/brick tiles ($M = 7.2$ vs. 6.0).

The number of targets found before any revisitation is a crude measure of how close a participant came to completing a trial perfectly, so we developed a more sophisticated metric that incorporated the proximity of a given participant to each target up until the first revisitation. Targets that had been found by this point scored zero, and the score for each remaining target was the closest distance (x) to the target so far (Figure 8). A repeated measures ANOVA showed a significant interaction between FOV and scene fidelity for the perfect search ($F(1, 36) = 4.40$, $p < .05$). Participants in the wide FOV/sports hall condition came closest to a perfect search with a mean score of 1.18m ($SD = 3.07$), whereas the wide FOV/brick tiles group were the furthest away ($M = 5.4$m, $SD = 9.7$). The other groups were normal FOV/brick tiles ($M = 3.0$m, $SD = 7.5$) and normal FOV/sports hall condition ($M = 4.7$m $SD = 8.0$).

The first behavioural measure was the amount of time participants spent performing different types of action. For each trial we calculated the percentage of time participants spent (i) planning where to go at the start of a trial), (ii) travelling around the VE, (iii) stationary between targets and decoys, and (iv) checking the target and decoy boxes. These data were then analyzed using repeated measures ANOVA's. Participants who used a wide FOV spent significantly less of their time planning ($F(1, 36) = 8.10$, $p < .05$), stationary ($F(1, 36) = 4.47$, $p < .05$), and checking the boxes ($F(1, 36)$

$= 8.71$, $p < .05$), than participants who used a normal FOV. It follows that participants who used a wide FOV spent significantly more of their time travelling around than participants who used a normal FOV ($F(1, 36) = 5.88$, $p < .05$). (Figure 6).
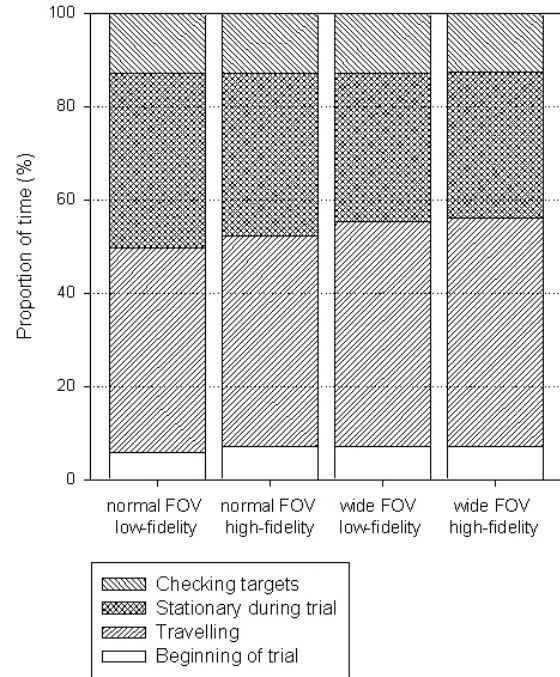
Figure 6: *Proportion of time spent on various actions during a trial*

The second behaviour measure was participants overall search strategy. The paths participants followed were classified by dividing the VE into four quadrants and noting the order in which these were visited (Figure 7).
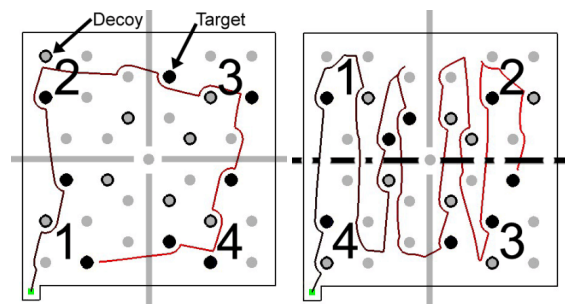
Figure 7: *Typical perimeter search (left) and lawnmower search (right)*

A perimeter search visited quadrants in the order 1-2-3-4-1 (clockwise search) or 1-4-3-2-1 (anticlockwise). A Lawn-

mower search involved a sequence of passes that crossed the VE's centreline, progressing along the centreline from one side of the VE to the other. Two searches in the tiled-texture/wide-FOV condition could not be unambiguously classified as perimeter or lawnmower, and so were termed as 'other'. Approximately half of the participants in every group completed the task by using the perimeter strategy, while the other half used the lawnmower strategy. These are the same strategies as were used by participants who performed the same task in a previous real-world study [LR03]

To explain the difficulties that participants encountered, three categories of errors where identified. This was accomplished by dividing a plan of the navigable environment into sectors using a Delaunay triangulation algorithm that used the centre of the cylinders, with additional points on the boundary walls and the four corners, as node points. For each trial, targets that were found after one or more targets or decoys had been revisited were classified into one of the three categories, depending on the participant's travelled path and its relationship to the neglected target. A *miss* was recorded if the participant had previously touched the cylinder on which the target's box was located (Figure 8a). *Local neglect* was recorded if the participant had previously travelled through any of the Delaunay triangles connected to the target's cylinder (Figure 8b). *Global* neglect was recorded for all other errors, indicating that the participant had not been in the target's immediate vicinity (Figure 8c).

There was one miss in each of the normal FOV conditions, six in the wide FOV/brick tiles, and four in the wide FOV/sports hall. The ratio of local:global neglect was similar in all four conditions but, in total, in the wide FOV/sports hall condition, participants made approximately half the number of errors than participants in both the normal FOV/sports hall and wide FOV/brick tiles conditions, and a third less than the normal FOV/brick tiles condition (Figure 9). Overall, 5% of the targets were missed, 27% were locally neglected, and 68% were globally neglected.

Further analysis of the path and direction of view of participants during trials discovered that out of the 31 locally neglected targets in the normal FOV conditions none actually came within the participants' FOV when they were in the targets immediate vicinity. This gives an explanation as to why so many targets fell into this category of error. It wasn't that participants neglected these targets, they simply didn't see them (Figure 10). However, out of the 32 locally neglected targets in the wide FOV conditions 17 were within the participants FOV, and all appeared in either the left or the right screen. Yet despite being visible and within close proximity to the participants the targets remained unsearched.

3. Discussion

This study investigated the effect of visual scene fidelity and FOV on participants' navigational ability in terms of task
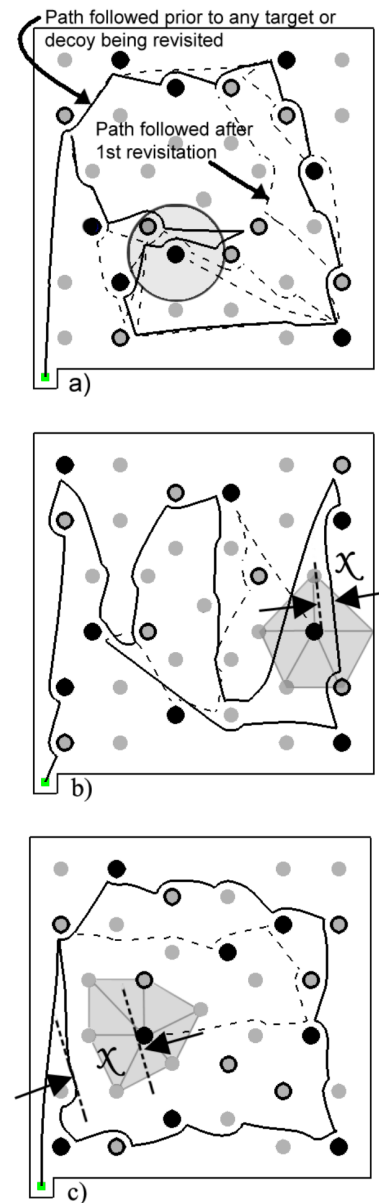
Figure 8: *Examples of the three types of error (a) miss, (b) local, and (c) global. 'x' shows the distance used in the 'perfect search' performance measure*

performance and behavioural metrics. With the combination of a wide FOV and a high fidelity (sports hall) scene participants were significantly closer to conducting "perfect searches" which is what is known to happen on the vast majority of occasions when people perform the same task in the real world.

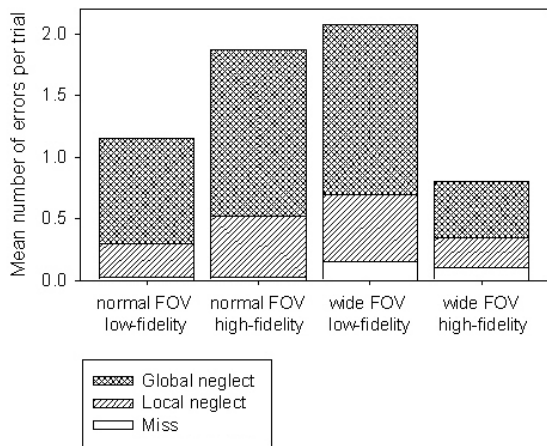Results from the behavioural metrics show significant dif-

Figure 9: *Type and mean number of errors for the four experimental conditions*
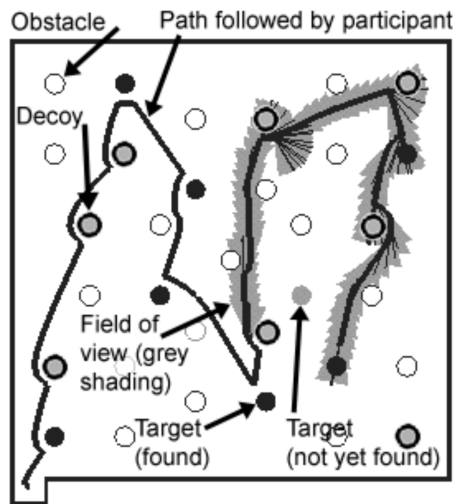

Figure 10: *Plan view showing a participant's path as they travelled near to a target. Despite the proximity of the participant to the target it did not come within the participants FOV and so remained unsearched*

ferences occurred between the two FOVs. With a wide FOV participants spent proportionally less time standing in one place at the beginning of a trial, between target and decoys, and checking them, than participants with a normal FOV. With a normal FOV, participants often stood in one place to look around and plan which box to search next, but a wide FOV allowed participants to assimilate the scene more easily simply because more of the VE was visible at any one time. It is reasonable to suggest that if more of the scene is outside the FOV then it is more demanding to keep track of the

locations of the possible targets and this then requires more looking around.

It should also be noted that looking around while travelling with a view-direction movement metaphor causes participants to veer off course. This metaphor is the one that is most commonly implemented in VEs, and was the one used in the present study. By contrast, observations from the real-world study in which participants walked around a physical version of the environment [LR03] showed that participants rarely stopped between targets and occasionally did not even stop to check boxes.

Participants searched the VE using two strategies, perimeter, or lawnmower. Approximately half of the participants in every group completed the task by using the perimeter strategy, while the other half used the lawnmower. The change in FOV did not change the frequency of one strategy over the other, unlike the real-world study [LR03] where the majority of participants who had a normal view of the environment searched the space with a perimeter strategy, while with a restricted FOV (20 degrees) the lawnmower strategy dominated.

The errors that prevented participants from conducting perfect searches were classified into three types, miss, local, and global neglect. Participants in the wide FOV/sports hall condition had far fewer instances of local and global neglect than the other three conditions. Errors occurred for three main reasons. First, both local and global neglect could occur when participants only searched some of the targets and decoys that existed in a tight cluster, whereas in the real-world participants tended to exhaustively search clusters when they were encountered. Second, entire clusters of targets were sometimes globally neglected in all the VE conditions, indicating that even with a wide FOV and a high fidelity scene, participants often had difficulty remembering where they had travelled. Lastly, in the two normal FOV conditions all of the locally neglected targets were outside the participant's FOV when they were within the targets immediate vicinity. Therefore it is unlikely that participants were even aware of the targets locations despite being in close proximity to them. In the wide FOV condition 53% of the locally neglected targets were within the participants FOV, and yet they were still not searched. Why where these targets neglected? It is possible that they where either not noticed or participants believed they had already been searched. If the former is true then it may have been caused by a difficulty participants had attending to the large physical area of the three-screen display. If the latter is true then one explanation might be that the participants where unable to assimilate the spatial information presented by the three separate screens into one useable representation of the sports hall. Even though more information was available they could not keep the positions of the targets updated with their own position and orientation as they travelled around the VE.

In conclusion, the present study shows that a shift towards

S. Lessels & R. A. Ruddle / Changes in Navigational Behaviour

real-world navigational behaviour occurs when a wide FOV is combined with a photorealistic scene. However, even with this combination of factors, performance remains substantially below that which occurs in the real world. In turn, this implicates the likely importance of more flexible forms of movement for efficient VE navigation.

4. Acknowledgements

This research was supported by Grant GR/R55818/01 from the UK Engineering and Physical Sciences Research Council.

References

[CTR02]   CZERWINSKI M., TAN D. S., ROBERTSON G. G.: Women take a wider view. In *Proceedings of CHI 2002* (2002), pp. 195–202. 2

[DS96]   DARKEN R. P., SIBERT J. L.: Wayfinding strategies and behaviors in large virtual worlds. In *Conference proceedings on Human factors in computing systems, CHI 1996* (1996), pp. 142–149. 1

[KKSB00]   KIRSCHEN M. P., KAHANA M. J., SEKULER R., BURACK B.: Optic flow helps humans learn to navigate through synthetic environments. *Perception 29*, 7 (2000), 801–818. 2

[KW96]   KLINE P. B., WITMER B. G.: Distance perception in virtual environments: effects of field of view and surface texture at near distances. In *Human Factors and Ergonomics Society 40th Annual Meeting* (1996), vol. 40, pp. 1112–1116. 2

[KWDT02]   KEARNS M. J., WARREN W. H., DUCHON A. P., TARR M. J.: Path integration from optic flow and body senses in a homing task. *Perception 31*, 3 (2002), 349–374. 2

[LR03]   LESSELS S., RUDDLE R. A.: Movement around real and virtual cluttered environments. Submitted for publication, 2003. 2, 5, 6, 7

[PMW97]   P.PÉRUCH, MAY M., WARTENBERG F.: Homing in virtual environments: Effects of field of view and path layout. *Perception 26*, 3 (1997), 301–311. 2

[RJ01]   RUDDLE R. A., JONES D. M.: Movement in cluttered virtual environments. *Presence: Teleoperators and Virtual Environments 10*, 5 (2001), 511–524. 1

[RMH99]   RICHARDSON A. E., MONTELLO D. R., HEGARTY M.: Spatial knowledge acquisition from maps and from navigation in real and virtual environments. *Memory and Cognition 27*, 4 (1999), 741–750. 1

[RPJ97]   RUDDLE R. A., PAYNE S. J., JONES D. M.: Navigating buildings on "desk-top" virtual environments: experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied 3*, 2 (1997), 143–159. 2

[SKD*99]   SINAI M. J., KREBS W. K., DARKEN R. P., ROWLAND J. H., MCCARLEY J. S.: Egocentric distance perception in a virtual environment using a perceptual matching task. In *Proceedings of the 43rd Annual Meeting Human Factors and Ergonomics Society* (1999), vol. 43, pp. 1256–1260. 2

[SM00]   STECK S. D., MALLOT H. A.: The role of global and local landmarks in virtual environment navigation. *Presence: Teleoperators and Virtual Environments 9*, 1 (2000), 69–83. 1

[SPRvdHB02]   SCHULTE-PELKUM J., RIECKE B. E., VON DER HEYDE M., BÜLTHOFF H. H.: Perceiving and controlling simulated egorotations from optic flow: influence of field of view (fov) and display devices on egomotion perception. In *Object Perception and Memory (OPAM) 2002* (2002). 2

[WBKP96]   WITMER B. G., BAILEY J. H., KNERR B. W., PARSONS K. C.: Virtual spaces and real world places: Transfer of route knowledge. *International Journal of Human-Computer Studies*, 45 (1996), 413–428. 1

[WHK98]   WALLER D., HUNT E., KNAPP D.: The transfer of spatial knowledge in virtual environment training. *Presence: Teleoperators and Virtual Environments 7*, 2 (1998), 129–143. 1, 2

[WK98]   WITMER B. G., KLINE P. B.: Judging perceived and traversed distance in virtual environments. *Presence: Teleoperators and Virtual Environments 7*, 2 (1998), 144–167. 2

© The Eurographics Association 2004.