White Rose
university consortium
Universities of Leeds, Sheffield & York

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

Roy A. Ruddle*
Stephen J. Payne
Dylan M. Jones
University of Wales
Cardiff CF1 3YG UK

Navigating Large-Scale ''Desk-Top'' Virtual Buildings:
Effects of Orientation Aids and Familiarity

Abstract

Two experiments investigated components of participants' spatial knowledge when they navigated large-scale "virtual buildings" using "desk-top" (i.e., nonimmersive) virtual environments (VEs). Experiment 1 showed that participants could estimate directions with reasonable accuracy when they traveled along paths that contained one or two turns (changes of direction), but participants' estimates were significantly less accurate when the paths contained three turns. In Experiment 2 participants repeatedly navigated two more complex virtual buildings, one with and the other without a compass. The accuracy of participants' route-finding and their direction and relative straight-line distance estimates improved with experience, but there were no significant differences between the two compass conditions. However, participants did develop significantly more accurate spatial knowledge as they became more familiar with navigating VEs in general.

1   Introduction

Experimental investigations and anecdotal evidence suggest that people frequently have difficulty navigating when they initially enter large-scale virtual environments (VEs; Darken and Sibert, 1996a,b; Henry, 1992; ''Research Directions in Virtual Environments,'' 1992; for a definition of large-scale space, see Weatherford, 1985). Factors that may contribute to this difficulty include people's lack of knowledge of their position, their orientation and a VE's structure, and a general lack of familiarity with using VEs, but the importance of each of these factors is currently under-researched.

The difficulties that people encounter while navigating have implications for the usability and effectiveness of VEs that are used in applications such as training, data visualization, virtual tourism, and virtual shopping. The primary purpose of some of these applications is to allow people to learn or investigate tasks that will later be performed in the real world (e.g., rescuing hostages, as discussed in Witmer et al., 1996). However, many people who explore VEs that are used for virtual tourism or virtual shopping, or are accessed via the World Wide Web will only experience the virtual version of the environment, even if a real-world ''replica'' exists. Therefore, research should address the navigation of VEs per se as well as the transfer of spatial knowledge learned in VEs to the real world.

*Electronic mail may be sent to ruddle@cardiff.ac.uk

This article presents the results of two experiments that investigated some of the factors that may influence people's ability to navigate VEs. The first experiment investigated the effects of changes of direction, field of view (FOV), and a compass on participants' ability to judge directions when they traveled along simple paths in virtual buildings. In the second experiment participants repeatedly navigated two, more complex virtual buildings, and these were used to investigate the effects of a compass. A secondary objective of Experiment 2 was to investigate whether changes occurred in the development of participants' spatial knowledge as they became more familiar using and navigating VEs in general. First we describe the background to these experiments that was provided by other VE navigation studies.

2   Navigation in VEs

Most VE navigation studies have been based on studies that investigated navigation in real-world situations (see Evans, 1980; Kitchin, 1994, for reviews). These studies often made distinctions between route- and survey-type spatial knowledge (for example, see Siegel and White, 1975; Wickens, 1992). Route knowledge is characterized by sequentially organized information about particular routes, whereas survey knowledge refers to the topographic properties of environments, for example, the positions of places relative to a fixed coordinate system and the straight-line distances between places (Thorndyke and Hayes-Roth, 1982). Similarities that have been found between spatial knowledge developed in VEs and in the real world (May, Péruch, and Savoyant, 1995; Tlauka and Wilson, 1996) suggest that the same distinctions are useful when considering VEs.

Some VE studies have investigated participants' ability to learn specific routes in virtual buildings (O'Neill, 1992; Witmer et al., 1996). Witmer et al. used a model of a real building. One group of their participants learned the route in the VE, made fewer errors as their training progressed, and successfully transferred their knowledge of the route when they were tested in the real building. However, these participants made significantly more errors during their training and during the test

than another group of participants who were trained and tested in the real building.

Participants in another study repeatedly navigated a 135-room virtual building and, after several hours, learned to route-find efficiently and also developed survey-type knowledge that was similar in accuracy to that of participants in an earlier study who navigated an equivalent real-world building (Ruddle, Payne, and Jones, 1997a; Thorndyke and Hayes-Roth, 1982). This suggests that the practical problem is not whether people can ever efficiently navigate large-scale VEs but how the development of people's spatial knowledge may be speeded up.

The VE displays used in the above studies provided a narrower FOV than people have in the real world. In one real-world study, participants learned the spatial layout of a room (small-scale space) less accurately when their FOV was restricted (Alfano and Michel, 1990) but in another study, which used a VE, no significant differences were found between the accuracy of participants' homing (direction) estimates when they used FOVs of 40°, 60°, and 80° (Péruch, May, and Wartenburg, 1997). Unpublished data from some of our VE studies showed that participants sometimes accidentally traveled past the locations for which they were searching when the locations lay just outside participants' FOV, but this accounted for less than 5% of their navigation errors.

The virtual building used by Witmer et al. contained a large amount of visual detail. Creating VEs with this detail is time consuming and expensive, but VEs that are of lower visual fidelity contain fewer visual cues and, therefore, potentially fewer landmarks. The route-finding by participants in the study by Ruddle et al. (1997a) was significantly more accurate when they navigated between locations in parts of a virtual building that contained landmarks at each corridor junction than when they navigated between locations in parts of the building that contained none of these landmarks. However, even with the landmarks, participants continued to have difficulty navigating from one location to another after spending several hours in the VE, and this suggests that landmarks alone are not sufficient to facilitate the rapid development of spatial knowledge (see also, Tlauka and Wilson, 1994).

Both landmark studies used local or "internal" (Evans

et al., 1984) landmarks that were only visible from within a restricted locality and provided only localized position and orientation information. Global or ''external'' landmarks, for example, a distant hill, the sun and the Pole (North) Star, are visible from far away and from many places. These landmarks provide people with information about their global (world-referenced) orientation but little information about their position. In a study in which participants navigated virtual seascapes that contained no barriers (e.g., walls) to movement, the addition of a virtual sun to a VE that contained landmarks seemed to help participants maintain their orientation and search more effectively for objects (Darken and Silbert, 1993). Similar information may be provided by displaying a compass within a VE so that it appears to be suspended just in front of participants, and the compass has the advantage of being visible no matter which direction participants are looking in.

Evidence from both this and another virtual seascape study (Darken and Sibert, 1996a,b) suggested that participants were able to return to their start position more quickly after searching for objects when a map was provided than when no supplementary aids were provided even though comparisons between these two particular conditions did not show statistically significant differences (R. P. Darken, personal communication, 1 November 1995). In a more recent seascape study (Ruddle, Payne, and Jones, 1997b) participants repeatedly searched for objects using aids that included a global map that showed the major topological features of the whole VE, a local map that only showed participants' immediate surroundings, but in greater detail, and both maps simultaneously (the L&G map). Participants learned the objects' positions significantly more quickly in each of the map conditions than when they navigated without any aids, and they learned quickest of all in the L&G map condition.

In summary, maps are an effective solution to the navigational difficulties that people encounter in VEs, but are of limited use in helping us understand the underlying nature of these difficulties. The following two experiments were principally designed to investigate two of these difficulties, the issues of orientation and a general lack of familiarity with VEs.
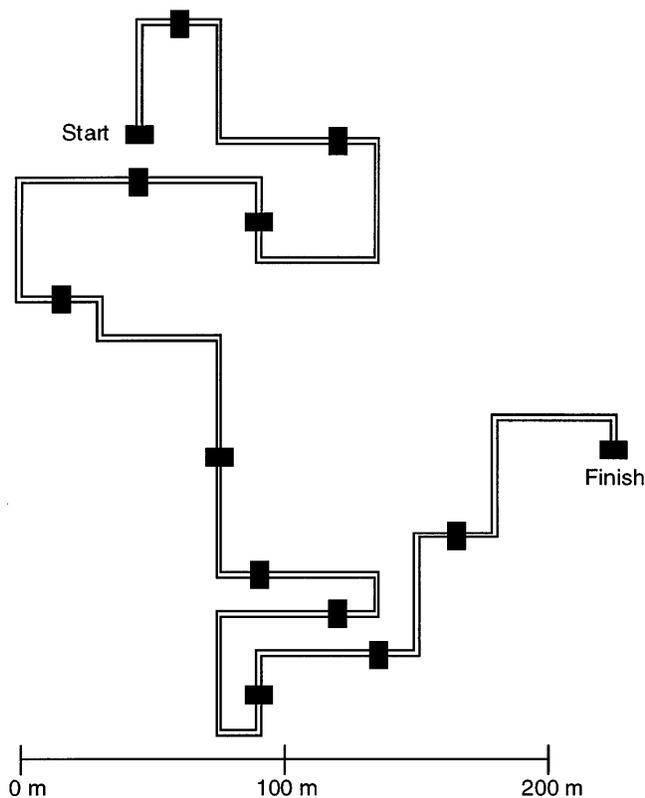
3   Experiment 1

The first experiment investigated how disoriented participants became when they traveled along simple paths in virtual buildings. Simple paths have been used to investigate aspects of spatial learning in a number of real world studies (e.g., Levine, Jankovic, and Palij, 1982; Levine, Marchon, and Hanley, 1984; Presson and Hazelrigg, 1984; Presson and Montello, 1994; Rossano and Warren, 1989). Each path in the virtual buildings led from one room to another and contained either one, two, or three 90° turns (changes of direction). In each room participants estimated the direction of the room they had come from. The experiment used a repeated measures design in which each participant made the estimates under four different conditions (with and without a compass, using 45° and 90° FOVs). The use of a within-participants design helped to overcome effects that were caused by differences in individuals' ability, and by any differences that may have been caused by various levels of experience in using computers, computer games, or a compass.

3.1 Method

3.1.1 Participants. A total of 16 participants (4 men and 12 women) took part in the experiment. They were all either undergraduates or graduates, who volunteered for the experiment and were paid an honorarium for their participation. Their ages ranged from 17 to 28 years ($M = 20.2$). The participants were divided into eight groups (two participants in each group) to counterbalance the order of the experimental conditions, the virtual buildings used for each condition, and the FOV participants used when they were familiarized with the VE controls.

3.1.2 Virtual Environment. The experiment was performed on a Silicon Graphics Crimson Reality Engine, running a C++ Performer application that we designed and programmed. A 21-in. monitor was used as a display and the application update rate was 20 Hz.
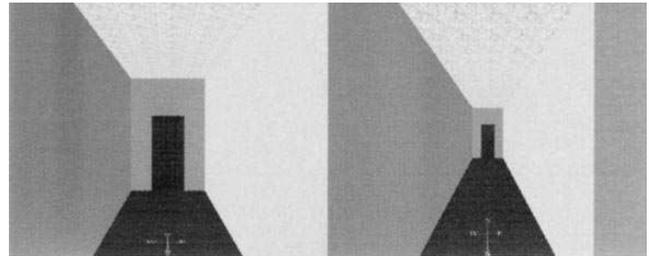
Six texture-mapped virtual buildings were created. Participants used one of these, a rectangular arrange-

Figure 1. *A plan view of one of the test buildings used in Experiment 1. The black rectangles indicate the positions of the 13 rooms.*

Figure 2. *A view inside one of the test buildings used in Experiment 1, showing the compass suspended in front of participants. The views are from the same position using a 45° FOV (left) and a 90° FOV (right).*

displayed, appeared to be suspended in front of participants. The compass rotated when participants changed their direction of view.

An interface, which allowed participants to travel in a straight line easily while simultaneously looking around, was provided by using the mouse and five keys on the keyboard. The mouse controlled the view direction in two ways.

- By moving the mouse from side to side, the view direction could be changed so that it panned through 180° (This was equivalent to participants turning their head from side to side.)
- By holding down the left or right mouse buttons, a full 360° rotation could be performed.

Four of the keys allowed participants to slow down, stop, speed up, and move at the maximum allowed speed (3 mph). The fifth changed participants' direction of movement to the current view direction. All participants mastered this interface without difficulty. At all times a green triangle, which projected at foot level, indicated the current direction of movement. Participants were prevented from walking through walls by a collision-detection algorithm, and doors opened automatically when approached.

3.1.3 Procedures. Participants were run through the experiment individually. First, a participant was familiarized with the VE controls and the procedure for estimating the directions using the rectangular practice building. Then the participant performed six tests in the

ment of corridors and two rooms, to familiarize themselves with the VE controls. The other five buildings were used to test participants' direction estimates. Each had a similar layout and one of these is shown in Figure 1. The buildings consisted of 13 rooms connected by paths made up of corridors that intersected at 90°. Four of the paths that connected the rooms contained one 90° turn, four contained two 90° turns, and the other four contained three 90° turns.

To define what was seen on the monitor, the application had to specify the height above the buildings' "floor" at which viewing took place (effectively a participant's virtual "eye" height) and the FOV to be used. Each participant's virtual eye height was set equal to their actual eye height, and participants navigated buildings with two different horizontal FOVs (45° and 90°). A typical view, using each of the FOVs, is shown in Figure 2. This figure also shows that the compass, when

virtual buildings. Participants performed the first three tests using one FOV and the remaining three tests using the other FOV. The first test in each block of three was treated as a practice and used the same building, but the other four (experimental) tests each used different buildings. For each participant the familiarization and six tests took a total of approximately 3 hr.

In each of the six tests participants started in the room at one end of the building and traveled from room to room, along the corridors. Each time they entered a room they pressed the ''y'' key to indicate their arrival and then the VE software moved them to the center of the room. Participants rotated their direction of view until they thought they were facing directly toward the room they had just come from and indicated this by pressing the ''y'' key, which caused the view direction to be recorded (the VE-orientation data). When participants had performed the direction estimate in the last room, the VE software exited.

3.2 Results

As expected, participants' direction estimates varied widely in accuracy. Participants' mean VE-orientation errors, averaged across the four experimental tests, ranged from 6° to 58°, and this result confirmed that our choice of a repeated-measures design was appropriate. The distribution of participants' VE-orientation errors was normalized using a logarithmic transformation and analyzed using a repeated measures analysis of variance (ANOVA). Figure 3 shows there was a main effect of number of turns on participants' mean VE-orientation errors, $F(2, 15) = 3.79$, $p < 0.05$. Planned contrasts showed that participants errors were significantly larger for rooms connected by three turns than for rooms connected by one turn, $F(1, 15) = 6.70$, $p < 0.05$, or two turns, $F(1, 15) = 4.44$, $p < 0.05$, but there was no significant difference in the errors for rooms connected by one and two turns, $F(1, 15) = 0.63$, $p > 0.05$. The same ANOVA showed that there were no significant difference between the 45° and 90° FOVs, $F(1, 15) = 0.29$, $p > 0.05$ ($M = 27$° vs. $M = 25$°), or between the compass and no compass conditions, $F(1, 15) = 2.54$, $p > 0.05$ ($M = 28$° vs. $M = 24$°).
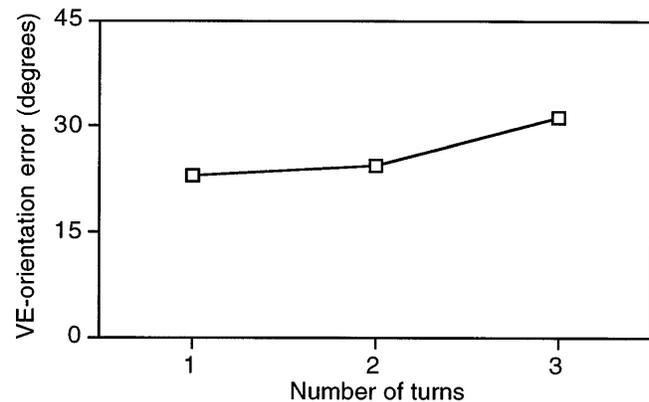

Figure 3. *Participants' mean VE-orientation estimate errors for paths that contained one, two, and three changes of direction.*

3.3 Discussion

The primary objective of this experiment was to determine how disoriented participants became when they followed simple paths in virtual buildings. The angular accuracy of participants' VE-orientation estimates may be put in perspective by comparing them with the accuracy of estimates made in other studies by participants who had learned the layout of real and virtual buildings. In a real-world study (Thorndyke and Hayes-Roth, 1982) the mean direction estimate error of participants who had worked in the building for between one and two years was 18°. In a study that recreated the Thorndyke and Hayes-Roth study in a virtual building participants' mean direction estimate error was 29° (Ruddle et al., 1997a). However, in this latter study most participants had not learned the shortest routes between all the test locations in the virtual building. (After spending an average of 4 hr in the VE, they still traveled an average of 10% farther than necessary.) Therefore, it is likely that the accuracy of their direction estimates would have further improved if they had navigated the building for a longer time.

The data from the study by Thorndyke and Hayes-Roth (1982) show how accurately people can estimate directions in familiar, real-world buildings. The data from the present study show that people can judge directions with an accuracy that approaches this level of accuracy if they follow paths in VEs that change direction once or twice, but the accuracy of their estimates

deteriorates significantly when they change direction more than twice. This finding suggests that people will have difficulty remembering the direction they have come from if they follow complex paths in VEs, even if these paths contain no places at which people must decide in which direction to travel.

There was no significant difference between the accuracy of participants' VE-orientation estimates when participants used 45° and 90° FOVs. The 45° FOV was approximately equal to the angle subtended by the 21-in. monitor when viewed from a normal viewing distance, whereas the 90° FOV distorted the image on the monitor but allowed participants to stand at the corridor junctions and look down both corridors simultaneously. This lack of an effect of FOV is in line with the findings of another study that found no effect of FOV when participants made homing estimates (Péruch et al., 1997).

The VE-orientation estimates that some participants made when the compass was not displayed had mean errors that were in excess of 50°. Despite becoming disoriented in that way, those participants seemed unable to use the compass to reduce the magnitude of their errors. Participants may have used the compass more effectively if they had been trained in its use and, in complex VEs, a compass might be used in different ways, for example, remembering the approximate direction of one location to another, and remembering the absolute position of locations in terms of compass bearings from a baseline reference point.

4   Experiment 2

The participants in Experiment 1 did not become completely disoriented when they followed simple paths in the virtual buildings. Instead, they made reasonably accurate estimates of direction, particularly when the paths only contained one or two changes of direction. However, most virtual buildings contain choices of routes, not just simple paths, and the decision points where these choices occur represent places where route-finding errors may be made.

Unpublished data from one of our earlier investiga-

tions using virtual buildings (Ruddle et al., 1997a) show that participants made as many route-finding errors at the first decision point of routes as at all other decision points combined. Therefore, one potential way of significantly improving participants' route-finding would be to help them make the correct choice at the first decision point, perhaps by supplying global orientation information. As has already been noted (see *Section 2* above), this information may be supplied in a number of ways, including the display of a virtual sun or a compass.

Experiment 2 had two objectives. The principal objective was to investigate the effects of a compass when participants repeatedly navigated two large-scale virtual buildings. These buildings (Building 1 and Building 2) were of similar complexity and their layouts are shown in Figure 4. The secondary objective was to investigate whether participants' spatial knowledge improved as a result of becoming more familiar with navigating VEs in general. To achieve the objectives, three central dimensions of participants' spatial knowledge were measured: (a) route-finding ability (distance travelled), (b) sense of straight-line distance (measured by calculating the Pearson correlation coefficient between a participant's estimated distances and the actual distances), and (c) direction estimate accuracy. The metrics and experimental design have been successfully used in other VE navigation studies (e.g., Ruddle, Payne, and Jones, 1997a; Ruddle et al., 1996), and some similar metrics have been used in other VE and real-world studies (e.g., Thorndyke and Hayes-Roth, 1982; Tlauka and Wilson, 1996; Wilson, Foreman, and Tlauka, in press).

4.1 Method

4.1.1 Participants. A total of 12 participants (7 men and 5 women) took part in the experiment. They were divided into four groups, which each contained at least one man and one woman. All were either undergraduates or graduates, who volunteered for the experiment, were different from the participants who took part in Experiment 1, and were paid for their participation. Their ages ranged from 19 to 29 years ($M = 21.4$).
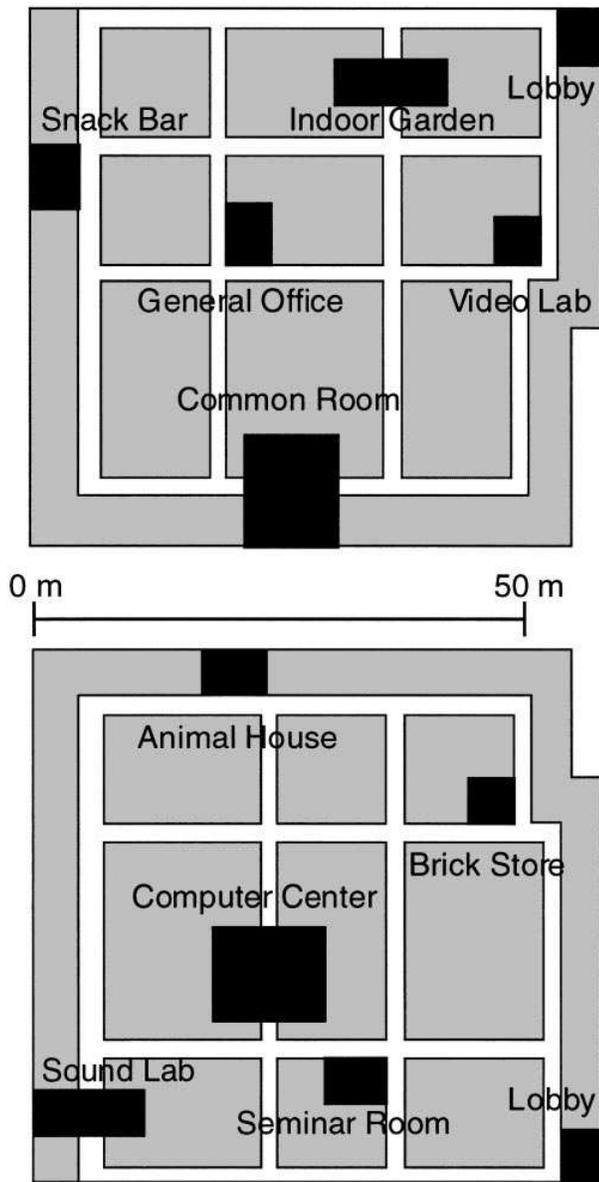
Figure 5. *A view inside Building 1. The view is from beside the video lab, looking toward the snack bar.*

Figure 4. *A plan view of the test buildings used in Experiment 2, showing the target locations (black), other rooms (gray), and corridors (white). Building 1 is the upper building.*

Participants in Group 1 navigated Building 1 eight times without a compass and then navigated Building 2 eight times with a compass. Participants in Group 2 navigated Building 1 with a compass and then navigated Building 2 without a compass. Participants in Groups 3 and 4 used the same building/ compass combinations as
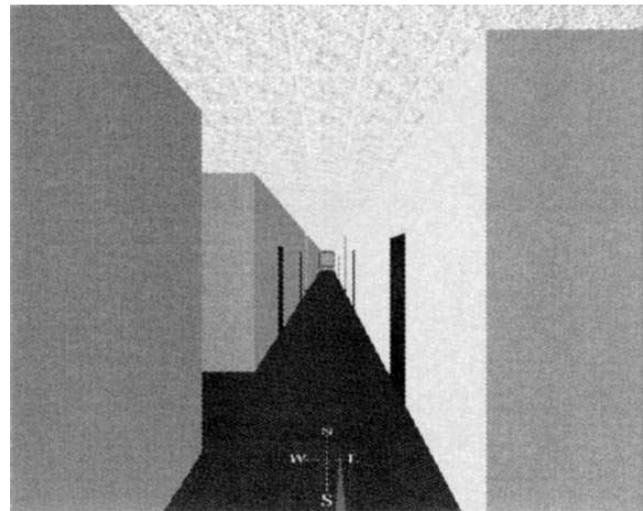
Groups 1 and 2, respectively, but navigated the buildings in the opposite order.

4.1.2 Virtual Environment. The experiment was performed using the same hardware, software application, and interface as Experiment 1, and a FOV of 90°. As in Experiment 1, the compass appeared to be suspended in front of participants and this, together with a typical view inside the buildings, is shown in Figure 5. Each building contained one lobby (vestibule) and five named rooms, which were filled with 3D models of characteristic furniture to enable their easy identification. The remainder of each building was divided into either 75 (Building 1) or 71 (Building 2) approximately equally sized empty rooms.

4.1.3 Procedure. Participants were run individually. First, a participant was familiarized with the VE controls using a simple practice building, which contained a figure-of-eight arrangement of corridors and two rooms, and then familiarized with the procedure for making the direction and distance judgments (see below). Then the participant navigated one test building eight times, and then navigated the other test building eight times, a process that took approximately 4.5 hr. To reduce fatigue,

participants came to our laboratory four times during one week and performed four navigation sessions each time.

The eight navigation sessions in each building were designed as virtual "days at the office" in which participants always started and finished in the lobby, and visited each of the other five named locations in an order which varied according to the session number. The days at the office were systematically structured and allowed our participants to experience a large proportion of the VE on several occasions, without being constrained to following specific routes. At the start of each session a message that was displayed on the screen named the first location that participants had to visit. When participants reached this location they pressed the "y" key, and this caused another message to be displayed, which named the next location to be visited, and so on. The messages were removed after a few seconds, but could be redisplayed at any time if the participant pressed the "h" key. At the beginning of each session a piece of paper was placed in front of the participant, which either said that the width of the common room (Building 1), or the computer center (Building 2) was 100 ft/ 30 m.

In Session 1 participants travelled to all locations by following a verbal description of the shortest route, which was spoken by the experimenter (e.g., "turn right out of the door, second left, and go through the door at the end"). In Session 2 participants followed verbal descriptions of the shortest route to the five named rooms but were told to find their own way back to the lobby, for which the following "2.5-min rule" applied.

If, after 2.5 min, a participant had not reached the lobby, the experimenter gave verbal instructions describing the shortest route to the lobby, which the participant then followed. However, if after 2.5 min, the participant was traveling directly towards the lobby, but had not yet arrived, they were allowed to continue unaided, but were given verbal instructions immediately if they deviated from the shortest route. No other form of feedback was given.

In the remaining six sessions (Sessions 3 to 8) participants navigated without help from the experimenter, but subject to the 2.5-min rule for each of the five rooms and the lobby. During all the sessions participants'

movements were recorded continuously for later analysis.

When participants arrived in each of the five named rooms in Sessions 5 and 8, they made estimates of direction and distance to the other four rooms. The direction (VE-orientation) estimates were made using the same procedure as Experiment 1. When a participant had made all four direction estimates, a Motif window was presented four times. Each time the participant entered an estimate for the straight-line distance from their current room to the named target room (the VE-Euclidean data; these distance estimates were termed "Euclidean" by Thorndyke and Hayes-Roth, 1982). All the estimates were from the center of the current location to the center of the target location and could be entered in meters or feet, according to the participant's preference.

After completing the test in the second building all participants answered a short written questionnaire that asked three questions: (a) Did you use the compass (yes/ no)? (b) Please list how you used the compass to find each room or to follow a route between particular rooms, and (c) How did your navigation differ when you did not have the compass?

4.2 Results

4.2.1 Data Analysis. Participants' route-finding ability in every unguided session was measured by computing the distance they traveled, in excess of the minimum possible distance, as a percentage of the minimum, the percentage extra distance travelled (PE-distance). Participants' appreciation of relative distance in the buildings was calculated by correlating their VE-Euclidean (straight-line) distance estimates with the corresponding actual distances. The distribution of this correlation was then normalized using Fisher's r-to-z transformation. Participants' direction estimate accuracy was determined by calculating the mean angular error of their VE-orientation estimates.

We wrote a second Performer application that overlaid the path participants traveled on to a plan view of the buildings. We used the application to determine where participants deviated from the shortest route when they

traveled to each location and, therefore, made their first route-finding error on each route.

4.2.2 Questionnaire Data. Nine of the participants indicated that they used the compass to help navigate the building. Five of these participants used the compass to help remember the positions of the five rooms in relation to the lobby, two tried to remember the rooms' and the lobby's position in terms of the four cardinal compass directions (North, South, East and West), one participant used the compass to help maintain the rooms' and lobby's general orientation and the other participant used the compass to help determine the rooms' and the lobby's positions relative to each other.

Participants indicated that they did not favor any particular strategy when they were not provided with the compass. Some participants learned the relative positions of the locations, others learned routes using landmarks provided by the buildings' structure (e.g., combinations of doors, and the zig-zags beside the video lab and the brick store), and other participants guessed which direction to travel in and hoped to find the locations by chance.

4.2.3 Navigation With and Without a Compass. The primary objective of this experiment was to investigate the effects of a compass on participants' spatial knowledge development. As in Experiment 1, participants varied considerably in their ability. Participants' PE-distance data was analyzed using a repeated measures ANOVA and used to compare their route-finding accuracy. Figure 6 shows that participants' route-finding accuracy improved significantly during the unguided sessions, $F(5, 11) = 16.25$, $p < 0.0001$. However, despite the different strategies used by participants in the two conditions (see Section 4.2.2 above) there was no significant difference between the compass and no-compass condition, $F(1, 11) = 0.10$, $p > 0.05$. Participants' means for the PE-distance data, averaged across the unguided sessions, were 79% and 83%, respectively.

The distribution of participants' VE-orientation data was normalized using a logarithmic transformation and
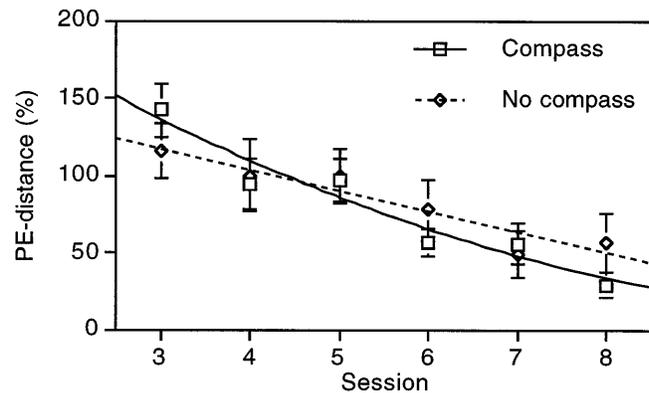
Figure 6. *Participants' mean percentage extra distance traveled (PE-distance) for the compass and no-compass conditions in Experiment 2. Error bars indicate standard error of the mean.*
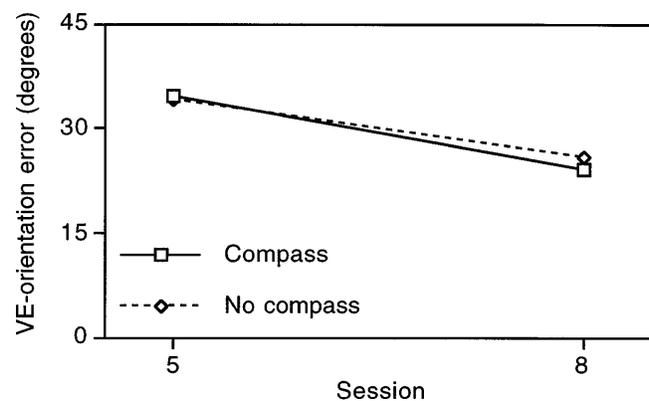
Figure 7. *Participants' mean VE-orientation estimate errors for the compass and no-compass conditions in Experiment 2.*

then analyzed using a repeated measures ANOVA. Figure 7 shows that participants made significantly more accurate estimates in Session 8 than in Session 5, $F(1, 11) = 23.34$, $p < 0.0005$, and planned contrasts showed that this difference was significant both when participants had a compass, $F(1, 11) = 8.30$, $p < 0.05$, and when they did not have a compass, $F(1, 11) = 5.03$, $p < 0.05$. However, there was no significant difference between the accuracy of estimates made in the compass and no-compass conditions, $F(1, 11) = 0.04$, $p > 0.05$.

Participants' sense of relative distance was also analyzed using a repeated measures ANOVA and showed a similar pattern of results to the VE-orientation data. Fig-
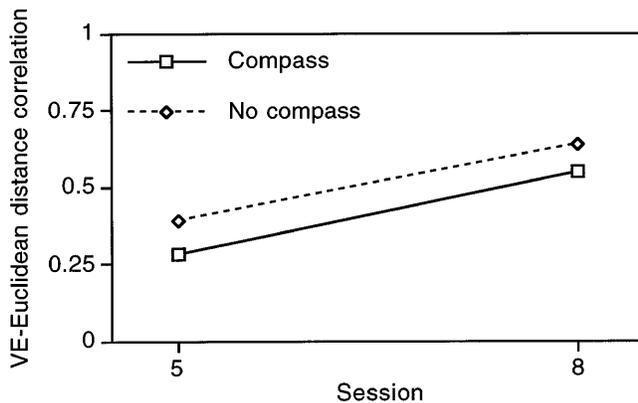
Figure 8. *Participants' mean VE-Euclidean distance correlations, transformed from participants' mean Fisher's z data, for the compass and no-compass conditions in Experiment 2.*

Figure 9. *Participants' mean percentage extra distance traveled (PE-distance) for the first and second buildings in Experiment 2. Error bars indicate standard error of the mean.*

ure 8 shows that participants' had a significantly more accurate sense of relative distance in Session 8 than in Session 5, $F(1, 11) = 38.82$, $p < 0.0001$, and planned contrasts showed that this difference was significant both when participants had a compass, $F(1, 11) = 10.14$, $p < 0.01$, and did not have a compass, $F(1, 11) = 10.98$, $p < 0.01$. Again, there was no significant difference between participants' sense of relative distance in the compass and no compass conditions, $F(1, 11) = 1.53$, $p > 0.05$.

4.2.4 The Effect of Familiarity with VEs on Navigation.

The secondary objective of Experiment 2 was to investigate variations in participants' spatial knowledge as they became more familiar with navigating VEs in general. We compared participants' PE-distance, VE-orientation, and VE-Euclidean data for the first building they navigated with the equivalent data for the second building. Participants in Groups 1 and 3 navigated the first building without a compass, whereas participants in Groups 2 and 4 navigated the second building without a compass.

A repeated measures ANOVA, illustrated in Figure 9, showed that participants' route-finding was more accurate in the second building than in the first building, $F(1, 11) = 64.26$, $p < 0.0001$. As in the above analyses, participants' VE-orientation data was normalized using a logarithmic transformation and analyzed using a re-
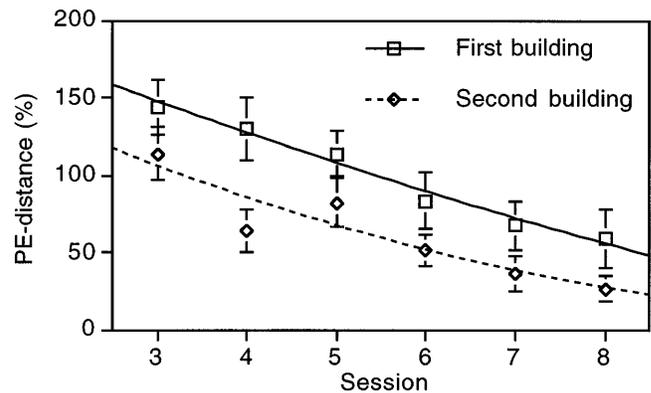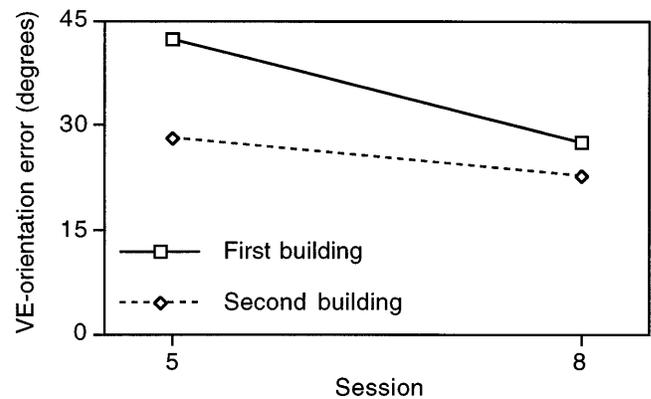

Figure 10. *Participants' mean VE-orientation estimate errors for the first and second buildings in Experiment 2.*

peated measures ANOVA. Figure 10 shows that participants' made significantly more accurate estimates in the second building than in first building, $F(1, 11) = 9.86$, $p < 0.0005$. Planned contrasts showed that this difference was significant for the estimates made in Session 5, $F(1, 11) = 12.88$, $p < 0.005$, but not for the estimates made in Session 8, $F(1, 11) = 2.58$, $p > 0.05$. Another ANOVA, illustrated in Figure 11, showed that participants' had a significantly more accurate sense of relative distance in the second building than in the first building, $F(1, 11) = 7.40$, $p < 0.05$, and planned contrasts showed that this difference was significant for both Ses-
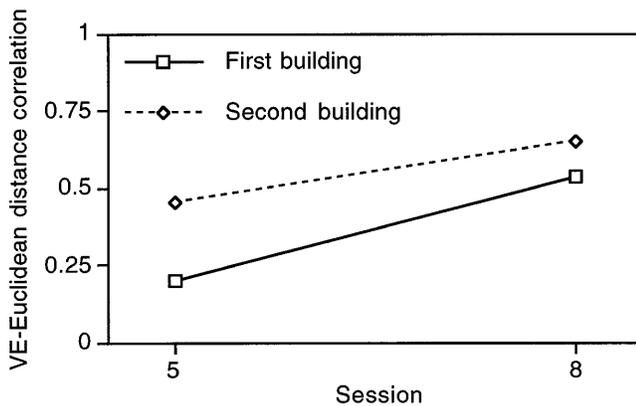
Figure 11. *Participants' mean VE-Euclidean distance correlations, transformed from participants' mean Fisher's z data, for the first and second buildings in Experiment 2.*

sion 5, $F(1, 11) = 7.63$, $p < 0.05$, and Session 8, $F(1, 11) = 15.47$, $p < 0.01$.

4.2.5 Route-finding Errors. The shortest route from one location to another contained from two to five decision points. These decision points occurred either where three or more corridor segments intersected, or where participants' current location had more than one exit (the common room, the computer center, the indoor garden, and the sound lab). The plan view software was used to compare the route participants travelled to each location with the shortest route to that location, and to classify any route-finding error in to one of four categories. These were: (a) direct (participants made no error and did not deviate from the shortest route), (b) miss (participants did not deviate from the shortest route until after they had travelled past the target location), (c) participants made their initial error at the route's first decision point, and (d) participants made their initial error at a subsequent decision point.

Figure 12 shows that participants' route-finding errors were similar, at least in broad terms, in the compass and no-compass conditions. Participants travelled directly to nearly 30% of the target locations in Session 3, and this proportion rose to 60% in Session 8. Participants missed their target on average of 3% of the routes. When participants made an error, it occurred as often at
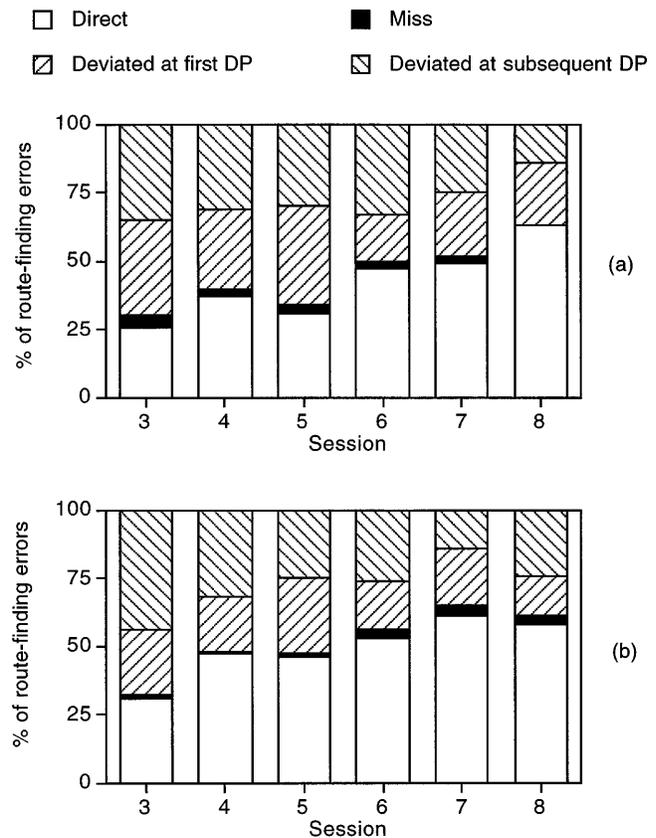

Figure 12. *Percentage of route finding errors in each error category for the compass (a) and the no compass (b) conditions in Experiment 2. DP = decision point.*

the first decision point as at all the other decision points combined.

4.3 Discussion

As expected and in keeping with other studies (e.g., Ruddle et al., 1997a; Tlauka and Wilson, 1994; Witmer et al., 1996) participants' mean route-finding ability improved with experience. However, even during the final (eighth) session, participants had considerable scope for improvement. The improvement in participants' route finding was mirrored by improvements in their survey knowledge, as measured by the accuracy of their VE-orientation and VE-Euclidean estimates.

The greatest differences occurred between the first building and the second building. Participants' route and survey knowledge was significantly more accurate in

the latter, particularly for the earlier sessions in each building (see Figures 9, 10, and 11). Part of this difference is likely to be due to the similarities between the two buildings (they were identical in size, contained the same number of decision points and were both based on a 4 × 4 matrix of corridors). Personal observations, supported by comments made by some participants, suggest that another factor was participants' increased familiarity with VEs in general when they navigated the second building. Initially, participants seemed to have little idea of the overall structure and size of the first building. Instead, they may have accepted that they were going to get lost and disoriented. By the end of Session 8 in the first building, participants had developed reasonably accurate spatial knowledge, even though, as has already been noted, there was still significant scope for improvement. This probably meant that participants quickly developed a general feeling for second building's structure and size, and this led to participants developing their spatial knowledge more quickly.

Surprisingly, especially given the differences in navigation strategy highlighted in the questionnaire data, there were no significant differences in participants' spatial knowledge in the compass/ no-compass condition. One possibility is that the buildings' structure allowed participants to find the target locations relatively easily by chance, and in more complex buildings participants would have had to navigate more accurately or risk never finding their target. However, this conclusion is not supported by the route-finding error data (see Figure 12), which showed that participants made a similar percentage of errors at the first decision point in each condition. If participants had used the compass to initially head in the correct direction then the proportion of errors made at the first decision point would have decreased.

Although the compass made no significant difference to the development of participants' spatial knowledge, the "comfort" provided by having global orientation information should not be ignored. One participant's answer to the third question was "When I did *not* have the compass, I was traveling blindly. There was no rhyme or reason as to where I went." The frustrations that people feel when they find computer application interfaces difficult or confusing are well known. Perhaps a

compass would make an important contribution to the confidence with which people use and navigate VEs, even if the compass had no significant effect on the accuracy of those people's navigation.

5   General Discussion

Participants in Experiment 1 were able to make reasonably accurate estimates of direction when they followed simple paths in virtual buildings. However, the introduction of complex routes (Experiment 2) led to participants having difficulty navigating, even after spending more than two hours in each virtual building. This suggests that the provision of global orientation information alone, via the display of a compass, is insufficient to help people quickly develop spatial knowledge. Different effects may have been found if other devices were used to provide this information, for example, a virtual sun, or color-coding the buildings' walls. Other effects may have been found if participants had been shown and allowed to practice a variety of search strategies that used the compass in ways such as memorizing the approximate direction of a baseline reference point such as the lobby.

A fundamental difference between the desktop VEs used in the present study and immersive VEs is that people physically turn to change their view direction in the latter. In some real-world studies that have used simple paths, participants estimated directions significantly more accurately if they physically rotated than if they imagined they had rotated (Presson and Montello, 1994; Rieser, 1989), and this may mean that people maintain their sense of global orientation more accurately in immersive VEs. However, in a study that used virtual buildings to compare participants' navigation when using desktop and immersive displays, no significant differences were found in route-finding or direction estimate accuracy (Ruddle et al., 1996) although it should be noted that the data were not conclusive and further investigation is required.

Perhaps both orientation and position information must be provided simultaneously to significantly affect

the accuracy of people's spatial knowledge. This could be achieved by combining an orientation aid with either local landmarks, or displaying people's current, momentary position using digital coordinates. Alternatively, both orientation and position information could be displayed on a map.

Route knowledge has been shown to develop more slowly in a VE than in an equivalent real environment (Witmer et al., 1996). People have a lifetime's experience of navigating in real-world situations, but most have only limited experience of using VEs, together with the resultant restricted FOV, lack of locomotion, often reduced visual fidelity, and lack of other modes of sensory feedback. Some studies have made preliminary investigations in to aspects of locomotion, including using a simulated walking interface (Slater, Usoh, and Steed, 1995) and physical movement using an omni-directional treadmill (Delaney, 1996; R. P. Darken, personal communication, 13 January 1997). Other studies are required to compare the route and survey knowledge that participants develop in VEs with the knowledge that they develop in real-world environments that contain the same amount of visual detail. Increased familiarity with VEs in general may allow people to adapt to the reduced amount of navigational information that is provided and lead to an increase in the rate at which spatial knowledge is developed. The first building vs. second building data in Experiment 2 provide initial support for this suggestion, but further investigations using buildings that are significantly different in size and structure are required before firm conclusions may be drawn.

References

Alfano, P. L., & Michel, G. F. (1990). Restricting the field of view: Perceptual and performance effects. *Perceptual and Motor Skills, 70,* 35–45.

Darken, R. P., & Sibert, J. L. (1993). A toolset for navigation in virtual environments. *Proceedings of ACM User Interface Software & Technology '93 (UIST '93),* 157–165.

Darken, R. P., & Sibert, J. L. (1996a). Navigating large virtual spaces. *International Journal of Human-Computer Interaction,* 49–71.

Darken, R. P., & Sibert, J. L. (1996b). Wayfinding strategies and behaviors in large virtual worlds. *Proceedings of Computer Human Interfaces Conference '96 (CHI 96),* 142–149.

Delaney, B. (1996). At last: An omni-directional treadmill. *CyberEdge Journal, 6.*

Evans, G. W. (1980). Environmental cognition. *Psychological Bulletin, 88,* 259–287.

Evans, G. W., Skorpanich, M. A., Gärling, T., Bryant, K. J., & Bresolin, B. (1984). The effects of pathway configuration, landmarks and stress on environmental cognition. *Journal of Environmental Psychology, 4,* 323–335.

Henry, D. (1992). Spatial perception in virtual environments: Evaluating an architectural application. Unpublished master's thesis, University of Washington. http://www.hitl.washington.edu/publications/henry/

Kitchin, R. M. (1994). Cognitive maps: What are they and why study them? *Journal of Environmental Psychology, 14,* 1–19.

Levine, M., Jankovic, I. N., & Palij, M. (1982). Principles of spatial problem solving. *Journal of Experimental Psychology: General, 111,* 157–175.

Levine, M., Marchon, I., & Hanley, G. (1984). The placement and misplacement of you-are-here maps. *Environment and Behavior, 16,* 139–157.

May, M., Péruch, P., & Savoyant, A. (1995). Navigating in a virtual environment with map-acquired knowledge: Encoding and alignment effects. *Ecological Psychology, 7,* 21–36.

O'Neill, M. J. (1992). Effects of familiarity and plan complexity on wayfinding in simulated buildings. *Journal of Environmental Psychology, 12,* 319–327.

Péruch, P., May, M., & Wartenburg, F. (1997). Homing in virtual environments: Effects of field of view and path layout. *Perception, 26,* 301–312.

Presson, C. C., & Hazelrigg, M. D. (1984). Building spatial representations through primary and secondary learning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 10,* 716–722.

Presson, C. C., & Montello, D. R. (1994). Updating after rotational and translational body movements: Coordinate structure of perspective space. *Perception, 23,* 1447–1455.

Research directions in virtual environments. (1992). *Computer Graphics, 26,* 153–177.

Rieser, J. J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15,* 1157–1165.

Rossano, M. J., & Warren, D. H. (1989). Misaligned maps lead to predictable errors. *Perception, 18,* 215–229.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1997a). Navigating buildings in ''desk-top'' virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied, 3,* 143–159.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1997b). *Navigating very-large-scale virtual environments: Search strategies and the effects of maps.* Manuscript submitted for publication.

Ruddle, R. A., Randall, S. J., Payne, S. J., & Jones, D. M. (1996). Navigation and spatial knowledge acquisition in large-scale virtual buildings: An experimental comparison of immersive and ''desk-top'' displays. *Proceedings of the 2nd International FIVE Conference,* 125–136.

Siegel, A. W., & White, S. H. (1975). The development of spatial representations of large scale environments. In H. W. Reese (Ed.), *Advances in child development and behavior* (pp. 9–55). New York: Academic Press.

Slater, M., Usoh, M., & Steed, A. (1995). Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Transactions on Computer-Human Interaction, 2,* 201–219.

Thorndyke, P. W., & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology, 14,* 560–589.

Tlauka, M., & Wilson, P. N. (1994). The effects of landmarks on route-learning in a computer-simulated environment. *Journal of Environmental Psychology, 14,* 305–313.

Tlauka, M., & Wilson, P. N. (1996). Orientation-free representations from navigation through a computer-simulated environment. *Environment and Behavior, 28,* 647–664.

Weatherford, D. L. (1985). Representing and manipulating spatial information from different environments: Models to neighborhoods. In R. Cohen (Ed.), *The development of spatial cognition* (pp. 41–70). New Jersey: Erlbaum.

Wickens, C. D. (1992). *Engineering Psychology and Human Performance* (2nd ed.). New York: HarperCollins.

Wilson, P. N., Foreman, N., & Tlauka, M. (in press). Transfer of spatial information from a virtual to a real environment. *Human Factors.*

Witmer, B. G., Bailey, J. H., Knerr, B. W., & Parsons, K. C. (1996). Virtual spaces and real-world places: Transfer of route knowledge. *International Journal of Human-Computer Studies, 45,* 413–428.