Journal of Experimental Psychology: Applied
1997, Vol. 3, No. 2, 143–159

Navigating Buildings in "Desk-Top" Virtual Environments: Experimental Investigations Using Extended Navigational Experience

Roy A. Ruddle, Stephen J. Payne, and Dylan M. Jones
University of Wales, Cardiff

Participants learned the layout of large-scale "virtual buildings" through extended navigational experience, using "desk-top" (i.e., nonimmersive) virtual environments (VEs). Experiment 1 recreated a study performed in a real building (P. W. Thorndyke & B. Hayes-Roth, 1982). After overcoming initial disorientation, participants ultimately developed near-perfect route-finding abilities. Their ability to judge directions and relative distances was similar to that found with the real building. Two further experiments investigated the effect of localized landmarks. Colored patterns had no effect on participants' route-finding accuracy. However, participants were more accurate in their route finding when familiar objects were used as landmarks than when no landmarks were used. The implications of the findings for the design of VEs are discussed.

Virtual environment (VE) systems are being investigated as an aid to training for real-world situations and for understanding complex data. A growing number of these VEs are "large scale" (e.g., Stansfield, Miner, Shawver, & Rogers, 1995), in that a user is unable to resolve the entire model from a single viewpoint (Weatherford, 1985). Critical to the effective use of such systems is the ability of the user to learn the spatial layout (i.e., develop a cognitive map) of the VE.

Empirical evidence suggests that users frequently have problems navigating VEs when supplementary aids (e.g., maps, artificial landmarks, etc.) are not provided (Darken & Sibert, 1996; Henry, 1992). Drivers of tele-operated vehicles (vehicles operated remotely by humans; see Sheridan, 1989, for a review of tele-operation) also experience navigational problems and have been unable to return to base once they wandered off a road, either intentionally or unintentionally (McGovern, 1991; personal communication, November 1994). Although these studies illustrate occurrences of navigational problems in VEs and tele-operation systems, none has quantified the extent to which navigation improves during extended use.

This article presents the results of three experiments that measured the quality of participants' spatial knowledge as they learned the layout of large-scale "desk-top" VEs (i.e., displayed using a monitor) through controlled, repeated experience. The first experiment allows close comparisons with a similar earlier experiment conducted in a real-world setting (Thorndyke & Hayes-Roth, 1982). The second and third experiments investigate the navigational benefits produced by including two different types of artificial landmarks in the VEs.

Background

There have been a number of studies that investigated the development of spatial knowledge in real-world situations, for example, towns, campuses, and buildings (see Evans, 1980; Kit-

---

Roy A. Ruddle, Stephen J. Payne, and Dylan M. Jones, School of Psychology, University of Wales, Cardiff, Cardiff, Wales, United Kingdom.

Correspondence concerning this article should be addressed to Roy A. Ruddle, School of Psychology, University of Wales, P.O. Box 901, Cardiff CF1 3YG United Kingdom.

chin, 1994; and Siegel & White, 1975, for reviews). When learning the layout of a large-scale environment through navigational experience, a person's spatial knowledge undergoes qualitative as well as quantitative changes, and it has been suggested that the qualitative changes are characterized by a progression through three levels of knowledge: (a) landmark knowledge, (b) route knowledge, and (c) survey knowledge (Siegel & White, 1975; Wickens, 1992). Buildings typically contain large numbers of route-finding decision points in a relatively small area and, as such, are useful environments for investigating navigational performance and the formation of spatial knowledge.

In the Thorndyke and Hayes-Roth (T & HR; 1982) study, the accuracy of participants' spatial knowledge of a building was compared, given one of two different learning methods: (a) employees with various levels of in situ experience of the building (navigation participants), or (b) study of a floor plan (map participants). The navigation participants made estimates of direction (survey-type knowledge) and route distance (route knowledge), which were significantly more accurate than the estimates of the map participants. In addition, the most experienced group of navigation participants (i.e., those who had worked in the building for 12–24 months) made estimates of straight-line distance that were of similar accuracy to those made by the map participants. These data were used to suggest that a person who has learned the layout of an environment through navigation ultimately develops spatial knowledge that is as accurate as—and, in some cases, more accurate than—the spatial knowledge developed by a person who has learned from a map (a survey-type perspective).

It is not known how spatial knowledge develops during extended navigational experience of VEs. Witmer, Bailey, Knerr, and Parsons (1996) reported a study in which participants successfully learned specific routes using a high-visual-fidelity VE model of a real building. The participants were able to transfer this knowledge when it was tested in the real building, although they made significantly more route-finding errors than participants trained in the real building. Comparisons have also been made between the performance of participants who navigated a real building and that of participants who navigated a low-visual-fidelity virtual model of the same building displayed using either a desk-top VE or an immersive VE (Henry, 1992). In this latter study, there were no significant differences between the three groups in terms of participants' directional accuracy when pointing to unseen locations. However, this lack of an effect of display mode may be explained by the simplicity of the building (seven rooms), which allowed participants to attain near-perfect spatial knowledge. Wilson, Foreman, and Tlauka (1996; see also Wilson & Foreman, 1993) compared the spatial knowledge of participants who navigated either a three-story real building or a low-visual-fidelity VE model of the same building until they were "familiar" with the layout. Half of each group of participants estimated directions in the real building while the other half of each group performed the estimates using the VE. Participants who navigated the real building made significantly more accurate direction estimates than participants who navigated the virtual building.

These studies suggest that either spatial knowledge is developed more quickly in the real world than in an equivalent VE or the ultimate accuracy of spatial knowledge developed in a VE is lower than that developed in the real world. The differences that exist between virtual and real worlds may help explain the differences between spatial knowledge development in the two media.

How Do Virtual and Real Worlds Differ?

When viewing a large-scale world from their normal, within-environment perspective, people must move through the environment to obtain all the information required to develop their spatial knowledge. The process of developing spatial knowledge thus involves integrating the information contained in each visual scene with a range of viewpoint locations and directions, changes in which are controlled by eye, head, and body movements over time (Weatherford, 1985). In desk-top VEs, movements are controlled by an abstract interface (mouse, keyboard, etc.) and the information that may be acquired from any single view is affected by the field of view and the presence of and type of cues that facilitate position and direction judgments. These differences may mean that spatial knowledge formed in a VE is different from that formed in the real world or is formed at a different rate. Specific sources of difference are described below.

Eye, Head, and Body Movements

In desk-top VEs, users receive feedback on their rotational and translational movements, which respectively cause changes of direction and position, solely from visual changes in the scene displayed. No vestibular or kinesthetic feedback is provided when users change their view direction, because eye, head, and body rotations are simulated using an abstract interface (e.g., a mouse or keyboard). Visual continuity during these changes in view direction is achieved by constraining the rate at which the view direction is allowed to change; even with a graphics supercomputer, the equivalent of a glance over the shoulder takes 1–2 s. The process of glancing becomes more like an implicit instruction to "rotate until you are facing the intended direction and then rotate back"; this changes the work required to integrate the information gained during the rotation with the user's existing spatial knowledge. Translational movements through VEs are also typically controlled using an abstract interface, and, therefore, users experience no physical locomotion.

The relative contributions of physical movements and visual experience when developing spatial knowledge are not known. Blind people have been shown to develop fairly accurate route- and survey-type spatial knowledge in a real building, although their knowledge was significantly less accurate than that of normally sighted people (Rieser, Lockman, & Pick, 1980). Other studies have shown that people make significantly greater errors in directional judgments when imagining their body has been rotated (a form of "abstract interface") than when physically rotating their body. No significant directional judgment differences were found between imagined translations and physical translations (Presson & Montello, 1994; Rieser, 1989).

Field of View

Hardware, distortion, and cost limitations typically restrict the field of view in VEs to 60°–100° at best. Operating with a restricted field of view increases the angle to which, and the number of times, users must rotate their heads to notice what they are walking past. The lack of peripheral vision has been shown to be important when learning the spatial layout of a room (Alfano & Michel, 1990).

Landmarks and Nonvisual Cues

VEs may be created with sufficient detail to bring their visual fidelity close to that of the real world. Unfortunately, this requires considerable time and cost, so more often than not, fidelity is compromised and the VE contains less detail and, potentially, fewer landmark-type cues than the real world. Senses other than vision are usually excluded from VEs, although there are few technical barriers to the inclusion of sound. The relative importance, however, of specific types and modes of information remains to be investigated in VEs.

Measuring Spatial Knowledge in VEs

Measurements taken within a VE are likely to be a better indicator of users' abilities to utilize their spatial knowledge in the VE than measurements taken afterward or in another perspective (e.g., tests on paper). This is because, in the VE, users have access to all the information present in their normal operating environment and the measurement and operating perspectives are identical.

In the following experiments, three central dimensions of spatial knowledge were investigated: (a) route-finding ability (distance traveled and time taken), (b) relative distance (measured by calculating the Pearson correlation coefficient between a participant's estimated distances and the actual distances), and (c) direction estimates (the angular error when "pointing" from one location to another).

Experiment 1

The first experiment investigated the development of participants' spatial knowledge when they learned the layout of a large-scale virtual building through navigational experience. Participants' route-finding ability was measured longitudinally during 10 different journeys through the building and, during the 10th journey, other components of each participant's spatial knowledge were measured using distance and direction estimates. The experiment recreated the T & HR study, substituting a virtual building for the actual physical facility used by those investigators.

Method

Participants

All 24 participants were either undergraduates or graduates who volunteered for the experiment and were paid an honorarium for their participation. They were split into two groups that each contained 6 men and 6 women. One group learned the building layout from a floor plan (the map participants); the other group (the navigation participants) learned from controlled exposure to the VE.

Materials

VE. The experiment was performed on a Silicon Graphics Crimson Reality Engine, running a C++ Performer application that we designed and programmed. A 21-in. (approximately 53 cm) monitor was used as a display, and the application update rate was 20 Hz.

Dimensions taken from the drawing of the RAND building presented in the T&HR (1982, p. 566) study were used to construct the VE database (see Figure 1), with the nine named locations (e.g., common room) filled with 3-D models of characteristic furniture to enable their easy identification. The remainder of the model was split into 126 roughly equally sized empty rooms. The entire VE building was texture mapped to help improve the visual realism of the walls, floor, ceiling, doors, and furniture. For convenience in generating the furniture, the supply room was changed to a video laboratory and the cashier was changed to a sound laboratory. There were no windows, and, apart from the features caused by the building's shape, there were no landmarks (plants, pictures, etc.) anywhere, other than the furniture in the nine named locations.

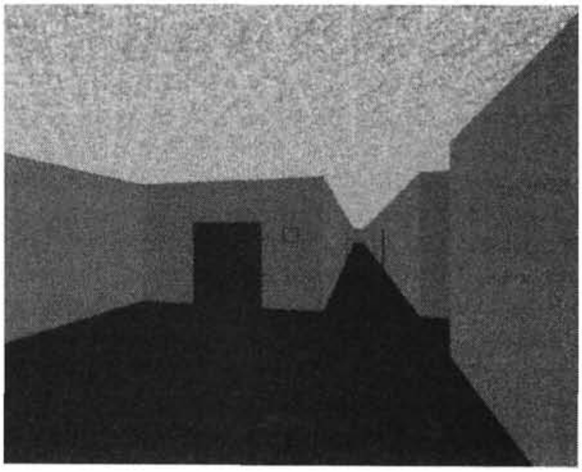

Figure 1. A view inside the VE used in Experiment 1. The view shows the open area outside the snack bar, and is looking toward the computer center.

To define what was seen on the monitor, the application had to specify the field of view to be used and the height above the building "floor" at which viewing took place (effectively, a participant's virtual eye height). Each participant was given the same horizontal field of view (90°) and eye height (1.70 m), and the vertical field of view (72°) reflected the monitor's aspect ratio. 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: (a) By moving the mouse from side to side, the view direction could be changed by ±45°, and (b) by holding down the left or right mouse buttons, a full 360° rotation could be performed. Four of the keys allowed the participant to slow down, stop, speed up, and move at the maximum allowed speed (3 mph, or 4.83 km/hr). The fifth changed the participant's 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.

Building floor plan (map). The building floor plan was presented on A3 paper (297 cm × 420 cm). The plan's scale was modified slightly (to 1 in.: 55 ft, or 1 cm: 660 m) for it to fit onto A3 paper.

Procedures

Participants were run individually and were told that the experiment was being performed to assess people's spatial knowledge in a virtual environment, given different types and amounts of learning experience. Each participant underwent two stages of training followed by a test. These are described in the following sections.

Training procedure. The first stage of training was the same for both groups. It was designed to allow the participants to become familiar with the VE controls using a simple virtual "practice building," which contained a figure-eight arrangement of corridors and two rooms. During the second stage of training, navigation participants learned the VE building's layout through repeated navigational experience and map participants learned from the floor plan.

Navigation participants learned the layout of the VE by undergoing nine training sessions, carried out approximately daily. These sessions were designed as virtual "days at the office," in which participants always started and finished in the same place, the East lobby (vestibule), and visited each of the other eight named locations in an order that varied according to the session number. The days at the office were systematically structured and allowed our navigation participants to experience a large proportion of the VE on several occasions, without having to "work" in the VE for as long as the T&HR navigation participants had in the RAND building (between 1 month and 2 years). During any particular training session (e.g., Session 1), each participant visited the named locations in the same order. Participants indicated they had reached each location by pressing the "y" key. This triggered the display of a message on the screen, which specified the name of the next location to be visited. The message was removed after a few seconds but could be redisplayed at any time if the participant pressed the "h" key.

In Sessions 1 and 2, navigation participants traveled to all except two of the named locations by following a verbal description of the shortest route, which was spoken by the experimenter (e.g., "turn left out of the door, left at the end, and go through the door"). In these sessions, the two exceptions were the East lobby (where the session had started) and either the snack bar or the common room (which were described as being opposite the common room or snack bar, depending on which had already been visited during that training session). For these two exceptions, participants tried to travel to the location as quickly as possible, but the following "5-min rule" applied.

If, after 5 min, a participant had not reached the new location, the experimenter gave verbal instructions that described the shortest route to the new location, which the participant then followed. However, if after 5 min the participants were traveling directly toward the target location but had not yet arrived, they were allowed to continue unaided but were given verbal instructions immediately if they deviated from the shortest route.

In the remaining seven sessions (Sessions 3 to 9), participants navigated without help from the experimenter but were subject to the 5-min rule. Participants' movements during training were recorded continuously for later analysis. They performed the test during a 10th session (see *Test procedure* below).

Map participants learned the building floor plan using the same procedure as in the T&HR experiment. Once they had memorized the plan (i.e., drawn it successfully, from memory, twice in succession), they performed the test (see *Test procedure* below). One participant was unable to redraw even the building outline after three attempts and was replaced in the experiment.

Test procedure. In keeping with the T&HR procedure, both groups performed the test inside the (virtual) building. All the participants first used the virtual practice building to familiarize themselves with the mechanism for answering the VE-orientation, simulated orientation, VE-Euclidean and VE-route questions (see below). For map participants, this process occurred immediately after they had satisfactorily memorized the map. For navigation participants, this familiarization occurred immediately before the final (i.e., 10th) session. Participants were then told that the distance from the center of the snack bar to the center of the common room, and the length of the computer center, were both 100 ft/30 m. Before the test, map participants had never been inside the virtual RAND building and navigation participants had never been shown the exterior of the building or a floor plan.

The sequence of events that took place during the test are shown in Figure 2. As in Training Sessions 3 to 9, navigation participants made their own way between all nine named locations and were subject to the 5-min rule. Map participants followed the shortest route between six of the named locations in the virtual building (the test locations: the video lab, computer center, conference room, East lobby, snack bar, and the South lobby), using verbal descriptions read aloud by the experimenter. At each test location, participants pressed the "y" key to indicate their arrival. They were then moved automatically to the center of the location, where there was a simple "compass," and automatically rotated to face the zero direction. The compass was a cross, aligned with each location's axes, with 0°, 90°, 180°, and 270° labeled. For each of seven target locations (Northwest lobby, sound lab, and the remaining five start points), participants answered three questions. First they rotated their viewpoint until they thought they were facing the required location and indicated this by pressing the "y" key, which caused the view direction to be recorded (the VE-orientation test). Then a Motif (UNIX) window was presented, and the participant entered estimates for the straight-line (Euclidean) distance and the shortest route through the hallways (the VE-Euclidean and VE-route distance data). All the estimates were from the center of the current test location to the center of the target locations and could be entered in meters or feet, according to the participant's preference. To prevent any ambiguity about the shortest route distance, the experimenter described the shortest route before that estimate was entered.

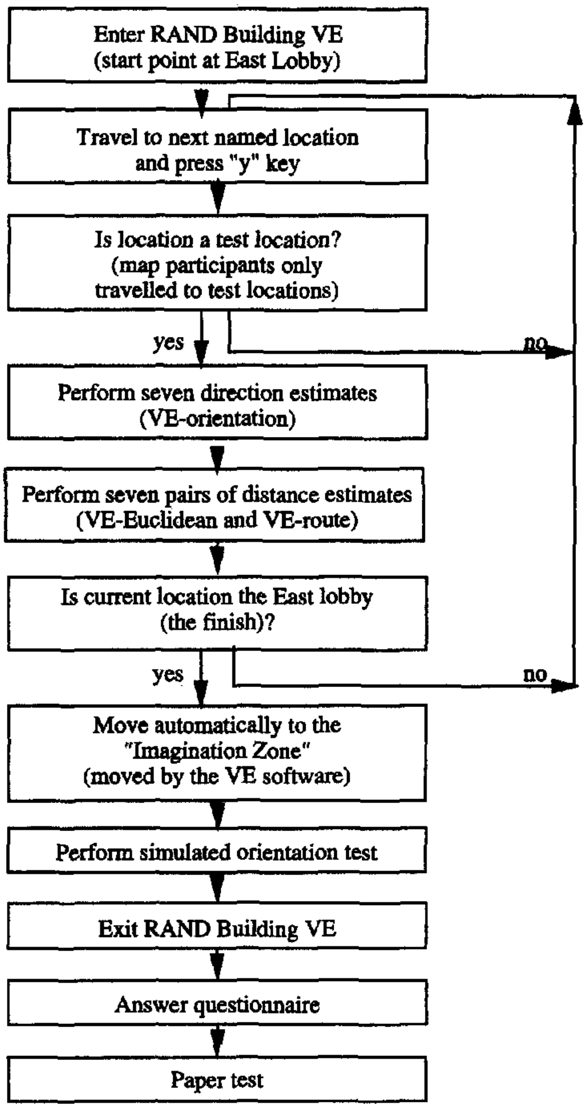

Figure 2. Test procedure used in Experiment 1. VE = virtual environment.

After returning to the East lobby (the sixth test location), the participant was moved automatically to a new room and performed the "simulated orientation" test (see Thorndyke & Hayes-Roth, 1982). All participants experienced great difficulty aligning themselves with this compass, and this was borne out in the wide variations in both groups' results. As a result, the simulated orientation data were not analyzed.

After the VE test, all participants answered a short written questionnaire that was designed to help us gain an insight into the approach they used to perform the test. Finally, participants were given a paper test that was essentially the same as the one used in the T&HR study and comprised 42 sheets of A4 paper (210 × 297 cm). On each sheet, there were two named crosses, one of which was circled, and, in the top left corner, the name of another location. Participants were told that the circled location was to be considered the start point and, using the second location as a reference, to mark a cross where they thought the center of the third location was. The 42 start–destination points were the same as in the VE tests. The first two locations were at an arbitrary orientation, but in all cases the scale was the same as the map used by the map participants and the third location always lay within the paper's bounds. Participants took the paper test away and completed it in their own time, and this may have compromised their responses. Therefore, data from the paper test have been omitted from the statistical analyses, but they have been included to allow a comparison with the paper angle (location) data from the T&HR study.

Results

Data Analysis

Participants' route-finding ability in every unguided session was measured by computing the distance they traveled to visit all nine target locations in excess of the minimum possible distance as a percentage of the minimum, the mean percentage extra distance traveled (MPED). The same calculation was also computed using time instead of distance. The time index had the advantage of accounting for periods when a participant was stationary but could not be used during the final session, as data for the distance and direction estimates were being gathered. Participant means for the two measures correlated very highly, $r = .96, p < .01$, so the distance metric was used for all further comparisons, as this could be calculated from eight, rather than seven, sessions.

Each participant's appreciation of relative distance in the VE was calculated by correlating their VE-Euclidean (straight-line) and VE-route distance estimates with the corresponding real-world distances. The distribution of these two correlations was normalized using Fisher's r-to-z transformation. We determined each participant's direction estimate accuracy by calculating the mean angular error of the VE-orientation and paper angle estimates.

Navigation Participants' Training

Navigation participants varied widely in the rate at which they learned the shortest routes between locations in the VE, but all the participants improved as the training sessions progressed (see Figure 3). Each participant's MPED, averaged across the eight unguided sessions, varied from 13% to 108%. Seven of 12 participants had near-perfect route knowledge by the time the test was performed, traveling less than 5% farther than necessary in the final session.

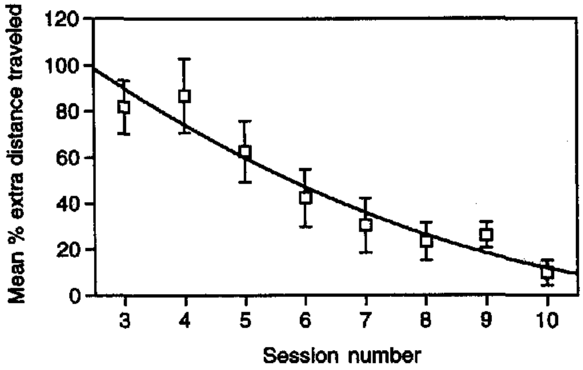

Figure 3. Navigation participants' mean percentage extra distance traveled in unguided sessions (Experiment 1). Error bars indicate standard error of the mean.

Estimates of Relative Distance

A multivariate analysis of variance showed that there was a significant main effect of training method on participants' VE-Euclidean and VE-route distance correlations, with map participants performing more accurately, $F(2, 21) = 4.46$, $p < .05$. Univariate analyses of variance (ANOVAs) were used to explore further the effect of training method on each type of distance correlation. The VE data in Figure 4 show that our map participants' correlations were significantly higher than those of navigation participants' for VE-Euclidean estimates, $F(1, 22) = 5.58, p < .05$,

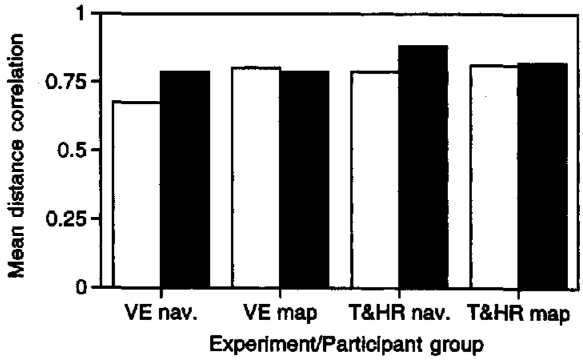

Figure 4. Mean distance correlations for Experiment 1 and Thorndyke and Hayes-Roth's (1982; T&HR) experiment. Open bars indicate Euclidean distance; solid bars indicate route distance. VE = virtual environment; nav. = navigation.

but our map participants' and navigation participants' VE-route correlations were similar, $F(1, 22) = 0.19, p > .05$.

Within-groups effects were analyzed using a repeated-measures ANOVA. Our navigation participants' VE-route correlations were significantly higher than their VE-Euclidean correlations, $F(1, 11) = 8.29, p < .05$, but there was no significant difference between our map participants' VE-route and VE-Euclidean correlations, $F(1, 11) = 0.56, p > .05$.

Estimates of Absolute Distance

Univariate ANOVAs, calculated using participants' absolute percentage distance estimate errors, showed that our map participants were significantly more accurate than our navigation participants at estimating the VE-Euclidean distances, $F(1, 22) = 6.08, p < .05$. However, the apparent difference for VE-route distance estimates was not significant, $F(1, 22) = 4.27, p > .05$ (see the VE data in Table 1). Navigation participants varied widely in accuracy and showed no consistent tendency to either under- or overestimate the VE-Euclidean or VE-route distances.

Direction Estimates

We analyzed the participants' mean VE-orientation errors using a univariate ANOVA. This showed that the apparent superiority of map participants' direction estimates, indicated by the VE data in Figure 5, was not significant, $F(1, 22) = 0.68, p > .05$.

When making the VE-orientation estimates, map participants first had to orient themselves with their memory of the map and then estimate the direction to each target location. Therefore, it was predicted that map participants would make more accurate VE-orientation estimates when they were in the lobbies, where their view down the corridors provided a visual frame of reference, than when they were in the rooms, where the door was always closed. Planned contrasts, using a repeated-measures ANOVA, showed that estimates from the lobbies ($M = 14°$) were significantly more accurate than those taken toward the lobbies ($M = 39°$), $F(1, 11) = 11.64, p < .01$, supporting this hypothesis. For comparison, the means for the navigation participants were 21° (from lobbies) and 18° (to lobbies).

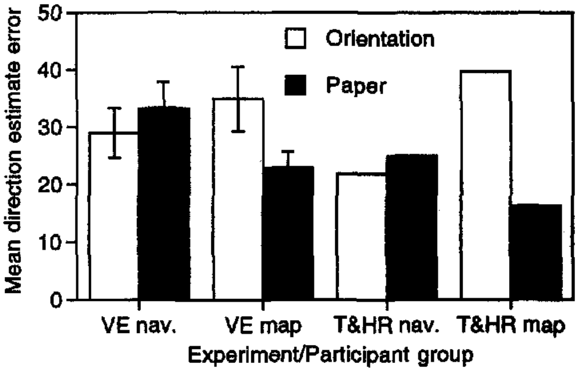

Figure 5. Mean direction estimate errors for Experiment 1 and Thorndyke and Hayes-Roth's (1982; T&HR) experiment. Error bars indicate standard error of the mean. VE = virtual environment; nav. = navigation.

Table 1
Means and Standard Errors for Absolute Percentage Error in Distance Estimates

Participant group	*M* for T&HR map	VE map *M*	VE map *SE*	*M* for T&HR navigation	VE navigation *M*	VE navigation *SE*
Euclidean	33	33	5	32	75	16
Route	36	35	5	26	63	12

Note. T&HR = Thorndyke and Hayes-Roth's (1982) study; VE = virtual environment.

Correlation Between Training and Distance and Direction Estimates

Navigation participants' MPED proved a reliable predictor of performance in the VE spatial knowledge tests. It correlated strongly with the three most important measures of spatial knowledge (i.e., those performed while inside the VE): mean VE-orientation angle error, $r = .92, p < .01$; VE-Euclidean distance estimate correlation, $r = -.82, p < .01$; and VE-route distance estimate correlation, $r = -.85, p < .01$.

Comparisons With T&HR

Detailed data from the T&HR study are not available to us, and, therefore, statistical comparisons with the present VE study are not possible. However, group means may be used to provide some comparisons between the two studies. In the comparison presented below, mean data for the T&HR navigation participants are averaged across the three levels of experience.

The Euclidean and route distance correlation data in Figure 4 show that our map participants' and the T&HR map participants' sense of relative distance was of similar accuracy. Our navigation participants had less accurate Euclidean and route distance correlations than T&HR navigation participants, but the difference was not large.

The data in Table 1 show that our map participants and the T&HR map participants also had similar mean absolute percentage errors for Euclidean and route distance estimates. However, our navigation participants varied considerably and, on average, performed more than twice as badly as the T&HR navigation participants.

The data in Figure 5 show an apparent difference in the accuracy of our map and the T&HR map participants' direction estimates. However, these differences were not particularly large in real terms (5° for VE-orientation and 7° for paper estimates). Our navigation participants made less accurate VE-orientation (7°) and paper direction estimates (8°) than the T&HR navigation participants, but again, the mean differences were not large in real terms.

To summarize, our navigation participants showed a substantial improvement during their training and had a significantly more accurate sense of relative VE-route distances than VE-Euclidean distances. Our navigation participants made less accurate direction and relative distance judgments than the T&HR navigation participants, but the differences were not large in real terms. Our map participants made significantly more accurate relative and absolute VE-Euclidean distance estimates than our navigation participants and also made distance and direction judgments of similar accuracy to the T&HR map participants. Our map participants used the VE interface less than our navigation participants did, because the former only used it to travel around the practice VE. If this had any effect, we expect that it would have reduced the accuracy of the map participants' estimates.

Discussion

The first question we sought to answer was whether VE users always remain somewhat disoriented when learning the layout of large-scale desk-top VEs through navigation and without navigational aids (maps, etc.). The answer appears to be no. After extended experience (approximately 4 hr, traveling about 15 km), an average navigation participant developed near-perfect route-finding ability and an ability to estimate relative distance (VE-Euclidean and VE-route distance correlations) and direction (VE-orientation angle error) that was similar in accuracy to that of the T&HR navigation participants, who had worked in the equivalent real building for 1–2 months or more. Participants developed this accuracy of spatial knowledge in a VE which contained only structural landmarks created by the building's layout (e.g., recognizable zigzags in corridors, dead-end corridors and distinctive groupings of doors) and did not have the wide range of additional cues (e.g., pictures, notices, plants) that are present in any real building. In fact, when moving through the VE, participants were never far from a structural landmark, but these were sometimes subtle in nature.

Like the T&HR navigation participants, our navigation participants had more accurate VE-route than VE-Euclidean distance correlations. However, further investigations are required to determine the longitudinal changes that occur in the accuracy of these data.

Despite being given a sense of scale, our navigation participants showed wide variability in their ability to estimate absolute distances, and this did not correlate with other measures of their spatial knowledge (route-finding ability, distance correlations and direction estimates). In addition, they did not demonstrate a specific tendency to either over- or underestimate the distances. Although we do not know how well they were able to estimate distances in the real world, their mean performance was substantially worse than that of the T&HR navigation participants. It may be that absolute distance estimation is inherently difficult in a VE. However, the difference could equally have been caused by the conflict between geometric field of view used to generate the VE display

(90°) and participants' physical field of view (approximately 40°; the angle subtended by the 21-in. [approx. 53 cm] monitor when seen from a normal viewing distance).

As expected given the replication of procedures, our map participants and the T&HR map participants developed a similar sense of both relative and absolute distance. These two groups' mean Euclidean and route distance correlations differed by less than 1 standard error (*SE*) of our map participants' judgments, and their mean absolute percentage errors for both Euclidean and route distance estimates differed by less than 0.2 *SE*. The accuracy of the direction estimates, which the two studies' map participants performed inside their respective environments, also differed by less than 1 *SE*.

Map participants' VE-orientation estimates were significantly more accurate when participants were provided with a visual frame of reference than when they had none. This suggests that map-learning users of VEs would benefit from a visual frame of reference, which allowed them to align their virtual selves with the environment. One method of providing such a frame of reference would be in the style of an aircraft head-up display.

Experiment 2

Experiment 1 demonstrated that people ultimately develop near-perfect route-finding ability and accurate survey-type spatial knowledge when they navigate a large-scale virtual building. However, this typically takes a long time. Empirical evidence suggests that we use a wide range of landmarks, for example, pictures, marks on walls, and structural features, to navigate real buildings, but the specific effects of landmarks are difficult to quantify. Heft (1979) showed that the inclusion of supplementary landmarks at path intersection points on woodland trails altered participants' navigational behavior but made no significant difference in the number of route-finding errors. Evans, Skorpanich, Gärling, Bryant, and Bresolin (1984) showed that landmarks aided the spatial relocation of photographs taken along a route (arranging the photographs spatially on a large piece of blank paper) but made no significant difference in the accuracy of route maps drawn after participants learned the layout of a simulated urban area. In VE studies, participants changed their navigational strategy when artificial landmarks were present and tried to use them as positional and orientational aids. However, no significant differences in navigation time, distance estimates, or direction estimates were found between landmark and nonlandmark conditions (Darken & Sibert, 1993, 1996, and personal communication, November 1995).

Experiment 2 investigated the effects of local or "internal" (Evans et al., 1984) landmarks and the rate of development of survey-type spatial knowledge when participants learned the layout of a second, large-scale virtual building through navigation. Participants' survey-type knowledge was measured using direction and distance estimates, performed during four different sessions.

Method

Participants

A total of 12 participants took part in the experiment. They were divided into two groups, which each contained 3 men and 3 women. 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. Both groups navigated the same VE, but each group used a different arrangement of landmarks. Figure 6 shows that the landmarks were positioned at each corridor decision point in one half of the building for the participants in Group 1 but were positioned at each corridor decision point in the other half of the building for the participants in Group 2.

Materials

The experiment was performed using the same hardware, software application, and interface as Experiment 1. Figure 6 shows that the virtual building was designed so that the two halves of the building were of similar complexity. The building contained one lobby and 10 named rooms (five in each half), which were filled with 3-D models of characteristic furniture to enable their easy identification. The remainder of the building was divided into 141 approximately equally sized empty rooms.

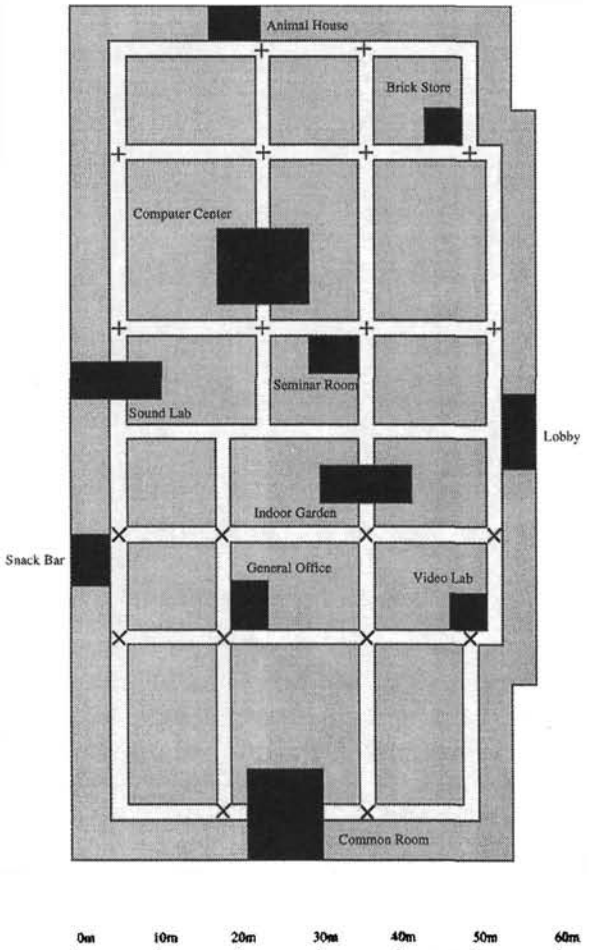

Figure 6. Floor plan of the virtual building used in Experiments 2 and 3. Landmark positions are marked with a + for Group 1 and an x for Group 2.

The landmarks took the form of solid, 1,000-mm high cuboids, had a 450 × 450 mm cross-section and were colored with "abstract paintings," which each had a different color scheme and pattern. Each column was symmetrical and, therefore, provided positional information but did not provide orientational information directly.

Procedure

Each participant was first familiarized with the VE controls, using the same practice building as Experiment 1. Then they learned the experiment VE's layout in nine sessions, which were structured as virtual days at the office, as in Experiment 1. The sessions were carried out in four blocks, which each lasted between 1.5 and 2 hr. Three sessions were completed during the first block and two sessions in the other blocks. Each participant completed all nine sessions during 1 week.

In each session, participants started in the lobby, visited the five named rooms in one half of the building, then visited the five named rooms in the other half of the building, and then returned to the lobby. The order in which the rooms were visited varied according to session number, and the sequence in which the building halves were visited was counterbalanced within each group.

In the first session, participants were guided round the building using the same type of verbal route descriptions as Experiment 1. In Session 2, participants were given explicit directions to find the named rooms, but they were not given directions when they had to return to the lobby at the end of the session. For the lobby, a 2.5-min rule applied, which used the same criterion as Experiment 1's 5-min rule. Experiment 2 used a 2.5-min time interval because the average distance between target locations was approximately half that in Experiment 1.

In the remaining seven sessions (Sessions 3 to 9) participants navigated without help from the experimenter, subject to the 2.5-min rule. In Sessions 3, 5, 7, and 9, participants made estimates of direction (VE-orientation) and straight-line (VE-Euclidean) distance from each of the target rooms to the other four target rooms in the half of the building occupied at the time. These estimates were made when participants arrived in each room and used the same interface (Motif window, etc.) as that used in Experiment 1, except there was no compass present in the rooms. At the end of Session 9, participants answered a short questionnaire that asked three questions: (a) "Did you notice the landmarks?" (b) "If so, did you use them?" and (c) "What were they?" Experiment 1's VE-route simulated-orientation and paper tests were not used in Experiment 2.

Results

Two MPED values were calculated for each participant, one each from the distances the participant traveled between rooms in the landmarked and nonlandmarked half of the building.

The experimental design was counterbalanced for any difference in the complexity of the two halves of the building, and, therefore, in the remainder of this section, the landmarked and nonlandmarked data are the combined data for both groups of participants.

The data in Figure 7 show that participants' route-finding ability, measured using the MPED metric, improved during the unguided sessions in both the landmarked and nonlandmarked halves of the virtual building. As in Experiment 1, participants varied considerably in their ability, with some still traveling more than twice the minimum distance during the final (ninth) session. Participants traveled longer distances between the landmarked locations than between the nonlandmarked locations ($M = 185\%$ vs. $M = 175\%$), but a repeated measures ANOVA showed that this difference was not significant, $F(1, 11) = 0.67, p > .05$.

As in Experiment 1, the distribution of participants' VE-Euclidean distance correlations was normalized using Fisher's r-to-z transformation and analyzed using a repeated measures ANOVA. The data in Figure 8 show the accuracy of participants' distance correlations improved with experience (session number), $F(3, 33) = 5.60$, $p < .005$, but there was no main effect of landmarks on distance correlations, $F(1, 11) = 0.18, p > .05$.

Participants' mean VE-orientation direction estimate error was also analyzed using a repeated measures ANOVA. The data in Figure 9 show the accuracy of participants' direction estimates also improved with experience, $F(3, 33) = 16.58, p < .0001$, but there was no main effect of landmarks on direction estimate accuracy, $F(1, 11) = 0.58$, $p > .05$.

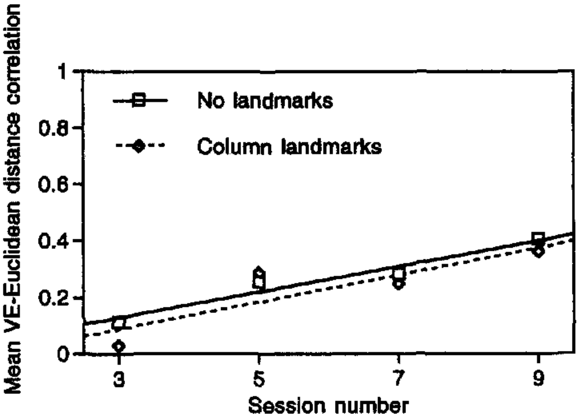

Figure 8. Mean VE-Euclidean distance correlations for Experiment 2 transformed from mean r-to-z values. VE = virtual environment.

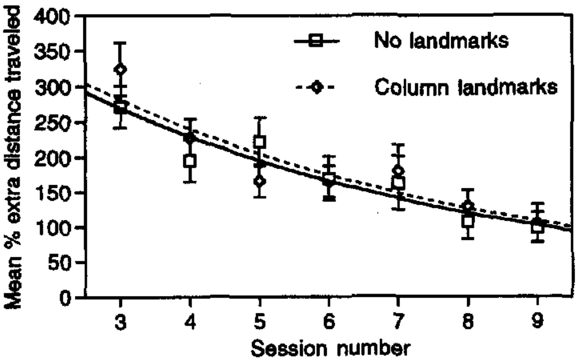

Figure 7. Participants' mean percentages for extra distance traveled in unguided sessions (Experiment 2). Error bars indicate standard error of the mean.

The questionnaire data showed that all the participants noticed the column landmarks and used them when navigating between the rooms. However, participants found it difficult to describe the landmarks, probably because they were in the form of abstract paintings rather than identifiable objects. Participants also indicated that the corridor zigzags by the brick store and the video lab, and the dead-end corridors, which lead to both doors of the common room, the indoor garden, the computer center, and the sound lab, were useful route-finding aids.

Discussion

As expected and in keeping with Experiment 1, participants' mean route-finding ability im-

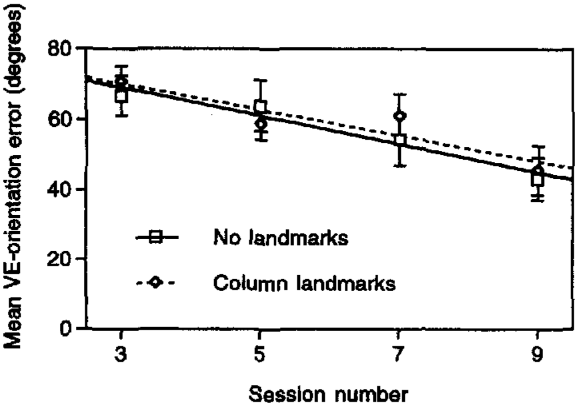

Figure 9. Mean VE-orientation direction estimate errors for Experiment 2. Error bars indicate standard error of the mean. VE = virtual environment.

proved with experience. However, even during the final (ninth) session, participants traveled an average of twice the minimum possible distance. Participants' MPEDs were much greater than in the equivalent sessions of Experiment 1, and we suggest this difference was caused by a difference in the two VEs' ease of navigation. One factor that may have contributed to this was the number of places participants had to decide in which direction to travel. The basic structure of the VE used in Experiment 1 consisted of a 3 × 5 matrix of corridors and an I-shaped wing, and had 13 corridor decision points (places where three or more corridor branches intersected). Experiment 2's VE was designed using a 4 × 7 corridor matrix and had 25 corridor decision points.

The distance correlation and direction estimate data in Figures 8 and 9 show that participants' survey-type knowledge increased during the unguided sessions. During the first of these sessions (Session 3), participants' estimates were only slightly more accurate than those which would be produced by random guesses (i.e., VE-Euclidean distance correlation = .00 and VE-orientation error = 90°). By the last session, participants had improved significantly, but their survey-type knowledge was still less well developed than that of participants who learned the layout of the virtual RAND building in Experiment 1.

The questionnaire data suggest that participants altered their route-finding strategy to try to make use of the column landmarks when traveling to rooms in the landmarked half of the building, but, despite having substantial scope for improvement, the landmarks made no significant difference to participants' route-finding ability or to their ability to judge directions or relative distances. This supports the findings of Darken and Sibert (1996) and Heft (1979), who also found that participants used landmarks when they were present, but the landmarks made no significant difference to participants' spatial knowledge development.

Experiment 3

The landmarks used in Experiment 2 were colored in the style of abstract paintings. Participants had difficulty describing these landmarks on the questionnaire and this may have led to difficulties when using them as navigational aids. Experiment 3 investigated the effects of using 3-D models of everyday objects as landmarks in the same large-scale virtual building as Experiment 2. We hypothesized that everyday objects were likely to be more distinctive and more memorable than the abstract paintings, and, as a result, participants might find it easier to form associations between the landmarks, the location of certain rooms, and the layout of the VE.

Method

Participants

A total of 12 participants took part in the experiment. They were divided into two groups, which used the same landmark positions as the corresponding groups in Experiment 2 (see Figure 6). Each group contained 3 men and 3 women. All the participants were either undergraduates or graduates who volunteered for the experiment, were different from the participants who took part in Experiments 1 and 2, and were paid for their participation.

Materials

We performed the experiment using the same hardware, software application, and interface as Experiments 1 and 2. The landmarks were solid cuboids, which were the same size as the cuboid landmarks used in Experiment 2 and were in the same positions. Each cuboid was gray and had a 3-D model of an everyday object placed on top. The objects were a car, a clock, a cup, a fork, a house, a cooking pot (a saucepan), a pair of glasses (spectacles), a traffic light, a toothbrush, and a truck. All the objects had a maximum dimension (width, breadth, or height) of 450 mm and had a familiarity rating of at least four out of five (Snodgrass & Vanderwart, 1980). The shape of the objects meant that each object gave orientational, as well as positional, information. For example, the fork was positioned near the general office for Group 1 and pointed toward the video lab.

Procedure

Experiment 3 used the same procedure as Experiment 2. A disk error caused some data for one participant to be lost during Sessions 3 and 6. We computed this participant's MPED using data from the other five unguided sessions, and the distance correlation and direction estimate data were omitted when participants' overall means were calculated for Session 3.

Results

The data in Figure 10 show that, as in Experiments 1 and 2, participants' mean route-finding ability improved during the unguided sessions (Session 8 used the same room order as in Session 4, and this order appears to have been easier than the three room orders used in the other sessions). Participants traveled shorter distances between the landmarked locations than between the nonlandmarked locations ($M = 138\%$ vs. $M = 178\%$) and a repeated measures ANOVA showed that this difference was significant, $F(1, 15) = 6.98, p < .05$.

The distribution of participants' VE-Euclidean distance correlations in Sessions 5, 7, and 9 was normalized using Fisher's r-to-z transformation, and we analyzed these data using a repeated measures ANOVA. The data in Figure 11 show the accuracy of participants' distance correlations improved with experience (session number), $F(2, 30) = 14.47, p < .0001$, but there was no main effect of landmarks on distance correlations, $F(1, 15) = 0.02, p > .05$.

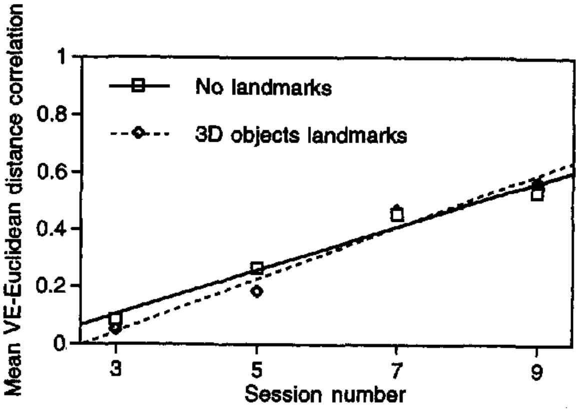

Figure 11. Mean VE-Euclidean distance correlations for Experiment 3, transformed from mean r-to-z values. VE = virtual environment.

Participants' mean VE-orientation direction estimate error for Sessions 5, 7, and 9 was also analyzed using a repeated measures ANOVA. Data in Figure 12 show that the accuracy of participants' direction estimates also improved with experience, $F(2, 30) = 24.36, p < .0001$, but there was no main effect of landmarks on direction estimate accuracy, $F(1, 15) = 0.10, p > .05$.

The questionnaire data showed that all the participants noticed the 3-D objects and used them when navigating between the rooms, typically by associating them with specific locations. As in Experiments 1 and 2, structural features were also used as route-finding aids.

Discussion

The development of participants' mean route-finding abilities, senses of straight-line distance, and direction estimate accuracies was similar to Experiments 1 and 2. As in Experiment 2, participants altered their route-finding strategy and used the landmarks when traveling to rooms

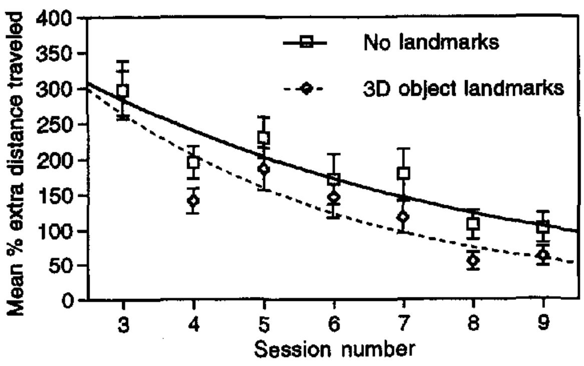

Figure 10. Participants' mean percentages for extra distance traveled in unguided sessions (Experiment 3). Error bars indicate standard error of the mean.

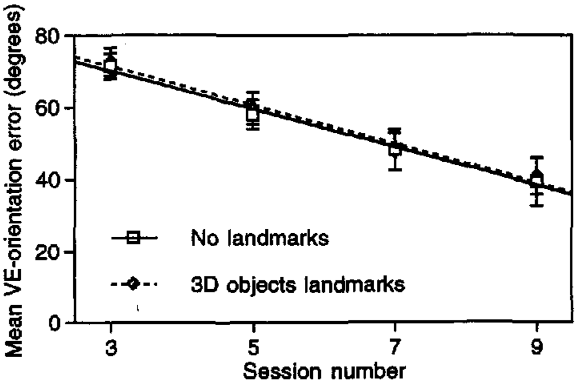

Figure 12. Mean VE-orientation direction estimate errors for Experiment 3. Error bars indicate standard error of the mean. VE = virtual environment.

in the landmarked half of the building. Unlike Experiment 2, participants traveled significantly shorter distances in the landmarked half of the building. However, to our surprise, there was no corresponding improvement in participants' direction and relative straight-line distance judgments. Further investigations are required to identify the reasons for this, but it is possible that the landmarks provided participants with confirmation of their approximate location (effectively enlarging the target area) rather than playing a significant role in the development of survey-type spatial knowledge.

Participants seemed to use the 3-D objects in two ways. The questionnaire data showed that all the participants formed associations between certain objects and rooms, and empirical evidence suggested that participants sometimes used these associations by first searching for an object and then searching a localized area for the target room. Participants may have formed stronger associations if the landmarks had been tied to the context of the environment (e.g., if a video cassette was used as a landmark near the video lab), but the effects of context-related landmarks remain to be investigated.

The disadvantage of forming object–room associations was that, when searching for rooms with two doors (the common room, the indoor garden, the computer center, and the sound lab), participants sometimes developed such a strong association with the object near one of the doors that they traveled past the object near the other door without realizing that they were close to their target room. If peripheral vision were supported, it may have reduced errors of this type. Informal observation suggests that the participants also used the 3-D objects by remembering a particular change of direction at an object that led them to a particular room (e.g., "turn right at the fork, and then the seminar room is on my left"). In this situation, participants sometimes changed direction but failed to find the target room, because they were approaching from a different direction to the direction they had been traveling in when they formed the association. After failing to find the target room in the expected location, participants then sometimes carried on, assuming their association was wrong, rather than backtracking to check their direction of approach.

General Discussion

The experiments described in this article demonstrate that users who navigate large-scale virtual buildings ultimately develop route-finding abilities and some survey-type spatial knowledge (direction judgments and relative straight-line distance judgments), which are as accurate as those abilities and spatial knowledge developed by people who work in real buildings.

Referring back to the three levels of spatial knowledge (landmark, route, and survey; Siegel & White, 1975; Wickens, 1992), our data illustrate the improvement of people's spatial knowledge when they repeatedly navigate a virtual building. In all three experiments, participants' route finding improved with experience and, in Experiment 1, ultimately became near perfect. Questionnaire data and informal observation suggest that participants used landmarks in two ways: (a) by forming associations between landmarks and the approximate position of target rooms, and (b) by using landmarks to trigger changes in direction of the routes between two locations. Participants' survey-type knowledge, as measured by their VE-Euclidean and VE-orientation data, also improved with experience, and this took place in parallel with the improvement in their route knowledge.

One explanation that is sometimes offered for the navigation and disorientation problems that occur when users initially travel through a VE is that the VEs do not contain many of the subtle cues (notices, marks on walls, etc.) that we use to navigate in the real world. This is true, but the creation of VE databases that contain a high level of visual detail is very time consuming and expensive, even when leading-edge software tools (Multigen, 3D Studio, etc.) and texture mapping techniques are used.

All the participants in Experiments 2 and 3 reported that they used the landmarks to help navigate the building, but only the memorable landmarks (3-D models of everyday objects) had a significant effect on the distance participants traveled. However, even during the final session with these landmarks, participants' route finding was far from perfect (they traveled an average of 61% farther than the minimum possible distance). Empirical evidence suggests that this extra distance traveled was usually the result of participants being partially lost or disoriented, rather than traveling to the location by an indirect route.

The implication for VE designers is that users are likely to be disoriented when they initially travel around a large-scale VE, even if it contains a high level of visual detail and, therefore, many landmark-type cues. In some applications, for example, virtual shopping or virtual tourism, users may be reluctant to use an application if they feel lost, feel disoriented, or have to spend a significant amount of time learning the VE's layout. Therefore, it is likely that alternative navigational aids must be developed.

One possibility is to supplement landmarks with orientation information; for example, a compass or a "virtual sun." The landmarks used in Experiment 3 gave positional information, but only subtle orientation information, the latter provided by the orientation of each 3-D object. Rieser (1989) and Presson and Montello (1994) showed that people made significantly larger errors in directional judgments when they imagined rotation movements of their body than when they imagined translationary movements. Therefore, the provision of a more obvious aid to orientation may lead to a significant reduction in route-finding errors, and a corresponding increase in the navigability and usability of large-scale VEs.

Another way to enhance a user's spatial knowledge when they initially use a VE is to provide "you are here" (YAH) maps or a small window within the VE that shows a miniature view of the whole VE (World in Miniature; Stoakley, Conway, & Pausch, 1995). Levine, Marchon, and Hanley (1984) demonstrated the superiority of an aligned YAH map over a contraligned map in a real building, where participants were significantly more successful at navigating between two locations and navigated significantly faster when using the aligned map. However, in another study which used a two-story building (Butler, Acquino, Hissong, & Scott, 1993), participants who had no prior knowledge of the building and who were given no navigational aids navigated between two locations more than twice as quickly as participants who were provided with a YAH map at start location. The participants who were not given any aids found the target locations by chance. Some of the extra time taken by the YAH participants was spent looking at and memorizing the map, but the size of the difference in the groups' mean route-finding times (approximately 3.5 min) suggests that the YAH participants also made route-finding errors. However, the YAH participants might have found the locations more quickly than the participants who had no aids, if the building had been more complex, and therefore, the locations were less likely to be found by chance. It should be noted that, unlike in the T&HR study, neither Levine et al. nor Butler et al. forced participants to memorize the YAH map perfectly, and this may account for some of the route-finding errors.

An alternative to memorizing maps is to provide maps which may be referred to at any time. This allows VE maps to be used in a similar manner to maps used as navigational aids in aircraft (e.g., Aretz, 1991) or when one walks through the countryside. Darken and Sibert (1996) investigated the effects of using no aids, a continuously displayed map, or landmarks when participants navigated a virtual seascape. Although a significant overall effect of aid type was found when participants revisited a location in the VE, and participants used different searching strategies with different types of aid, the differences between each pair of aids were not significant. Therefore, further research is required to investigate the effectiveness of maps for navigation in large-scale VEs and to optimize the designs of these maps.

References

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

Aretz, A. J. (1991). The design of electronic map displays. *Human Factors, 33,* 85–101.

Butler, D. L., Acquino, A. L., Hissong, A. A., & Scott, P. A. (1993). Wayfinding by newcomers in a complex building. *Human Factors, 35,* 159–173.

Darken, R. P., & Sibert, J. L. (1993). A toolset for navigation in virtual environment. *Proceedings of the ACM Symposium on User Interface Software & Technology '93* (*UIST '93,* pp. 157–165). New York: ACM.

Darken, R. P., & Sibert, J. L. (1996). Navigating large virtual spaces. *International Journal of Human–Computer Interaction, 8,* 49–71.

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.

Heft, H. (1979). The role of environmental features in route-learning: Two exploratory studies of wayfinding. *Environmental Psychology and Nonverbal Behavior, 3,* 172–185.

Henry, D. (1992). *Spatial perception in virtual environments: Evaluating an architectural application.* Unpublished master's thesis, University of Washington. Retrieved from the World Wide Web May 15, 1995. 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., Marchon, I., & Hanley, G. (1984). The placement and misplacement of you-are-here maps. *Environment and Behavior, 16,* 139–157.

McGovern, D. E. (1991). Experience and results in teleoperation of land vehicles. In S. Ellis (Ed.), *Pictorial communication in virtual and real environments* (pp. 182–195). London: Taylor & Francis.

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

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.

Rieser, J. J., Lockman, J. L., & Pick, H. L., Jr. (1980). The role of visual experience in knowledge of spatial layout. *Perception & Psychophysics, 28,* 185–190.

Sheridan, T. B. (1989). Telerobotics. *Automatica, 25,* 487–507.

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.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 174–215.

Stansfield, S., Miner, N., Shawver, D., & Rogers, D. (1995). An application of shared virtual reality to situational training. *Proceedings of the Virtual Reality Annual International Symposium '95 (VRAIS 95,* pp. 156–161). Los Alamitos, CA: IEEE.

Stoakley, R., Conway, M. J., & Pausch, R. (1995). Virtual reality on a WIM: Interactive worlds in miniature. *Proceedings of Computer Human Interfaces Conference '95 (CHI 95,* pp. 265–272). New York: ACM.

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

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). Hillsdale, NJ: Erlbaum.

Wickens, C. D. (1992). *Engineering psychology and human performance* (2nd ed.). New York: HarperCollins.

Wilson, P. N., & Foreman, N. (1993). Transfer of information from virtual to real space: Implications for people with physical disability. *Proceedings of the First Eurographics Workshop on Virtual Reality* (pp. 21–25). Aire-la-Ville, Switzerland: Eurographics Association.

Wilson, P. N., Foreman, N., & Tlauka, M. (1996). *Transfer of spatial information from a virtual to a real environment.* Manuscript submitted for publication.

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.

Received March 12, 1996
Revision received January 28, 1997
Accepted January 31, 1997 ■