Published article:

Ruddle, RA, Payne, SJ and Jones, DM (1999) *The effects of maps on navigation and search strategies in very-large-scale virtual environments.* Journal of Experimental Psychology: Applied, 5 (1). 54 - 75.

http://dx.doi.org/10.1037/1076-898X.5.1.54

The Effects of Maps on Navigation and Search Strategies in Very-Large-Scale Virtual Environments

Roy A. Ruddle
Stephen J. Payne
Dylan M. Jones
Cardiff University

[contact Ruddle for figures]

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

Abstract

Participants used maps and other navigational aids to search desk-top (non-immersive) virtual environments (VEs) for objects that were small and not visible on a global map which showed the whole of a VE and its major topological features. Overall, participants searched most efficiently when they simultaneously used both the global map and a local map which showed their immediate surroundings and the objects' position. However, after repeated searching the global map on its own became equally effective. When participants used the local map on its own their spatial knowledge developed in a manner which was previously associated with learning from a within-environment perspective, rather than a survey perspective. Implications for the use of maps as aids for VE navigation are discussed.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

The Effects of Maps on Navigation and Search Strategies in Very-Large-Scale Virtual Environments:

Although people tend to be very disoriented when they initially navigate a large-scale virtual environment (VE), they ultimately develop accurate route- and survey-type spatial knowledge (Ruddle, Payne, & Jones, 1997). However, the process of developing this knowledge typically takes a long time. One reason for this is that people must explore the environment and integrate information seen from many different positions because they can not resolve all the detail necessary for efficient navigation from a single position if a "human's-eye" perspective is used (Siegel, 1981; Weatherford, 1985).

Some large-scale environments, for example, a section of a city, an office containing partitions, or the VEs used by Ruddle et al. (1997), may be resolved from a single position if a plan view perspective is used. The detail of these environments may be displayed in a single view on a map and this facilitates rapid learning. The size of some other environments, compared with the relative size of important features or objects within the environments, means that these environments may not be resolved from a single position in any perspective and, therefore, may be termed very-large-scale. Maps used as navigational aids for these environments may be presented at a variety of different scales, each showing different amounts of the environment and levels of detail. For example, one map may use a small scale to show the whole environment in a relatively small amount of detail (a global map), whereas another map may use a larger scale yet be manageable because it scrolls as people travel so that only the part of the environment surrounding them is shown, but in greater detail (a local map). Both types of map may take the form of dynamic YAH (you-are-here) maps on which a person's position is regularly updated, and which are displayed continuously to allow navigation information to be looked up on demand, rather than having to be memorized.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

Examples of very-large-scale VEs include some of the VEs which are used in the construction industry (see "Engineering and Construction," 1997), military applications such as the control of uninhabited air vehicles (UAVs; teleoperated aircraft), and some applications for the tourism industry (for example, see "Virtual Tourism: Paris", 1996) and the World-Wide Web. One challenge for designers is to make very-large-scale VEs navigable so that people may find their way around without undue difficulty. This article presents the results of two experiments which investigated the effectiveness of local maps, global maps and other aids when participants navigated very-large-scale VEs, and the effects that these aids had on participants' search strategies. The experiments were performed using desk-top VEs (i.e., a monitor display). Other types of VE display include those that are visually immersive (see "Headmounted Displays," 1997) and those that are spatially immersive (theatre-based systems; see "Spatially Immersive Display Systems," 1998).

Background

Despite the difficulties that people have experienced while navigating VEs and the importance of navigation (see Encarnação, Foley, Bryson, Feiner, & Gershon, 1994; "Research Directions in Virtual Environments," 1992), the effects of different types of aids remain under-researched. Studies which have investigated aids for use in VEs have either used "virtual seascapes" which consisted of large areas of ocean and a few islands (Darken & Sibert, 1993, 1996a; see also, Darken & Sibert, 1996b), or "virtual buildings" (Ruddle et al., 1997; Ruddle, Payne, & Jones, 1998; see also, Tlauka & Wilson, 1994). One difference between the two types of environment was that the walls in the virtual buildings were "solid" barriers which constrained people's movement to defined routes, whereas the seascapes contained no such barriers and, therefore, allowed unconstrained movement. In VEs that allow unconstrained movement, for example, those used to control UAVs, people can travel directly to any location by using their survey knowledge (e.g., compass bearing and straight-line distance), but in constrained-movement VEs people also require knowledge of possible routes.

In the seascape studies, participants first searched for particular objects that were distributed throughout the seascapes (naive or uninformed searches) and then returned to their start point (primed or informed searches). The first study (Darken & Sibert, 1993) performed baseline investigations of the effects of a variety of aids. The second study (Darken & Sibert, 1996a) investigated the effects of four display conditions in greater detail, and these were: (a) no aids, (b) a forward-up, dynamic YAH map that did not show the objects' position, (c) a grid and five tall landmarks that were visible from any position in the VE, and (d) the map, grid and landmarks.

The data from the second study did not provide clear evidence of the superiority of any one particular aid. There was no effect of display condition on the time that participants took to perform the uninformed searches and, although there was a significant main effect of display condition on participants' informed search times, none of the differences was significant between pairs of conditions (R. P. Darken, personal communication, November 1, 1995). Reasons for this include: (a) Participants only had to search for a single point (the start point), and (b) The scale of the VEs and the visibility of the start point was such that participants could find the start quickly by chance.

The two studies did provide useful information about the strategies which participants used to search VEs and the effects on those strategies of different navigational aids. Overall, the most common search strategy was edge following (searching around the island and VE boundaries), but this meant that participants sometimes traveled more than once over the same part of the VE (an inefficient search) and this strategy also led to difficulties when participants were asked to return to their start point as they had little knowledge of their global position or orientation. Participants who used a north-up (fixed orientation), dynamic YAH map seemed to travel shorter distances when performing the uninformed search than participants who used a forward-up (track-up) map, which rotated so that it was always aligned with participants' view of a VE, but participants who used the forward-up map seemed to have more accurate knowledge of their orientation.

Participants sometimes used landmarks, when provided, to divide the seascapes into segments that could then be searched separately, and sometimes used landmark clusters, for example, those arranged in an "L" shape, to provide direction information. Ruddle et al. (1997) showed that landmarks aid navigation in large-scale virtual buildings. In that study, participants formed associations between certain landmarks and the locations they were searching for, and used these associations to either localize the search area, or remember that a particular change of direction at a landmark led to a particular location.

A variety of sources of error affect the accuracy with which people can judge directions in VEs (e.g., see Ellis, Smith, Grunwald, & McGreevy, 1991). When a virtual sun (without shadows) was added to seascapes that contained landmarks, participants seemed to maintain more accurate knowledge of their global direction (Darken & Sibert, 1993). Although the virtual sun obeyed the principle of pictorial realism (see below), it was only visible when it lay inside participants' field of view (FOV). A graphical compass also provides global orientation information, obeys the principle of pictorial realism and may be displayed so that it appears to be suspended in front of a person, giving it the advantage of being visible all of the time. Ruddle et al. (1998) investigated the effect of a compass when people repeatedly navigated virtual buildings. Most participants changed their strategy to use the compass but it had no effect on the accuracy of their route or survey knowledge. However, there may have been a greater likelihood of an effect occurring if the building had been larger or more complex, or participants had been taught how to effectively use the compass.

The most efficient way to exhaustively search the whole of one of the seascapes during an uninformed search was to use a lawnmower pattern. This strategy seemed to be used most often when participants were given position information via continuously-updated Cartesian XY coordinates (Darken & Sibert, 1993). With other aids, for example, a map, this strategy required participants to overlay VEs with an imagined (conceptual) mown grass pattern.

In addition to these directly relevant studies, the design of VE navigational aids may be guided by fundamental principles which have been developed for aids used in aircraft, including the principles of pictorial realism, the moving part, visual momentum, and optimum scaling (Aretz, 1991; Roscoe, 1968; Roscoe, Corl, & Jensen, 1981; Woods, 1984). These are discussed in the following section.

Design Principles for Navigational Aids

Pictorial Realism

The principle of pictorial realism states that displays are more effective when information is encoded graphically and can be readily identified with its meaning in the real world (Roscoe, 1968). Therefore, in VEs it may be more effective to display people's position graphically on a plan view outline of an environment than by displaying numerical X and Y coordinates.

The Moving Part and Visual Momentum

The principle of the moving part was developed during aircraft control studies and concerns the question of whether an environment should appear to move relative to the aircraft, or the aircraft relative to the environment (Johnson & Roscoe, 1972; Roscoe, 1968). Related to this principle are the differences between ego- and world-centered reference frames (respectively ERFs and WRFs). Aids that use an ERF present information with respect to pilots' current (momentary) position and orientation and an example is a forward-up map. Aids that use a WRF present information with respect to a more stabilized, world coordinate system and include north-up maps.

As with the VE study by Darken and Sibert (1993), studies using aircraft simulators have shown that both forward-up and north-up maps offer advantages over each other for certain types of tasks (for example, see Barfield, Rosenberg, & Furness III, 1995; Wickens, Liang, Prevett, & Olmos, 1996). Participants in the latter study made significantly faster and significantly more accurate judgments of the relative direction of terrain features when flying with a forward-up map, but made significantly faster judgments of features' absolute bearings and misplaced fewer features when drawing a sketch map after flying with a north-up map. Barfield et al. used three-dimensional, perspective views in which either the environment rotated (an ERF) or a symbol marking the aircraft's position rotated (a WRF). Participants locked-on to and acquired targets significantly more quickly, and flew a significantly more accurate flight path when the ERF view was used, but positioned the targets significantly more accurately on a plan view map when the WRF was used. Hooper and Coury (1994) found similar effects when they investigated the use of north-up and forward-up graphical displays for providing orientation information in submarines. Studies performed using simple paths and in buildings have consistently shown that aligned (forward-up) maps are more effective navigational aids than non-aligned maps (Levine, Marchon, & Hanley, 1984; Levine, Jankovic, & Palij, 1982; Presson & Hazelrigg, 1984; Rossano & Warren, 1989), but participants in these studies only performed ego-centred tasks.

Wickens & Prevett (1995) suggested that a compromise between an ERF and a WRF display was to provide an ERF-based, 3-D view of an environment in which the viewpoint was tethered to pilots' position and orientation by a constant offset. This was then investigated experimentally and, although participants made more accurate world-referenced position judgements while they were flying with the tethered display than with a forward-up map, participants in both conditions had great difficulty constructing a map after the flight had finished.

An alternative to using a tethered display is to use the principle of visual momentum (Woods, 1984) to provide a visual link between the WRF used by a north-up map and the ERF perspective of pilots' view from inside aircraft. Aretz (1991) showed that the addition of visual momentum in the form of a graphical wedge to a north-up map provided an optimal solution which combined the advantages offered by the north-up map for world-centered tasks with the advantages offered by the forward-up map for ego-centered tasks. In that study, participants flew a low-level helicopter simulation and, at specific times were asked to compare their marked position on the map with their view out of the cockpit and say whether the position marked on the map was correct. Participants' reaction times were significantly quicker when using either the forward-up map, or the visual momentum map than when using the north-up map. After completing each flight, participants were asked to sketch a map of the environment, and they correctly positioned significantly more landmarks when using either the north-up map, or visual momentum map than when using the forward-up map.

A note of caution comes from Williams, Hutchinson and Wickens (1996), who compared participants' spatial knowledge of an environment when they flew a route in an aircraft simulator either with, or without, a map. Participants who flew with the map drew significantly more accurate sketch maps after the flight but got just as lost as the other participants when asked to fly back to their start point, and this suggests that the map participants were unable to apply their superior spatial knowledge to a route-finding situation.

Optimum Scaling

The principle of optimum scaling is concerned with the scale of a display. It is dependent on the type of environment and tasks that are to be performed and, therefore, should be investigated experimentally (Roscoe, 1968). With maps, the optimum scale is a compromise between a small-scale map that shows all the places a person might go and a larger-scale map that helps the person improve their navigational precision. The effects of scale on VE navigational aids are discussed in the following section.

Navigational Aids for Very-Large-Scale VEs

The above VE and aircraft studies show that dynamic YAH maps are one candidate for effective navigational aids in very-large-scale VEs. These maps obey the principle of pictorial realism, because they graphically show people's position in environments, and may be designed using visual momentum to incorporate the advantages of both maps that have a constant (north-up) orientation and forward-up maps that are always aligned with environments.

The amount and detail of a VE that may be displayed in a single view of a map depends on the map's scale. A small-scale (global) map may show the whole of a VE and its major features in a single view. This type of map theoretically gives people sufficient information to conduct an efficient uninformed search because they can see their global position at all times and, therefore, should be able to search all of an environment by traveling over each part once. A global map may also allow people to perform an efficient informed search if they can remember objects' approximate global positions, because they will be able to travel directly to them. One disadvantage of a global map is that, due to it's small scale, it is unlikely to show the position, shape or identity of small objects and features. Therefore, people can only see objects and features by looking in their direction while in their immediate vicinity.

As the scale of a map increases, more detail becomes visible but people must integrate information that is learned from many different positions to develop accurate spatial knowledge. From a practical point of view, a local map shows the immediate surroundings of an object, but does not directly show its global position or the relative position of two objects if they lie more than a certain distance apart. People also have to integrate information from many different positions when they use a human's-eye perspective to learn the layout of large-scale environments by navigation, and this substantially increases the time that they require to develop accurate spatial knowledge when compared with learning from a single, plan-view perspective (a map; see Ruddle et al., 1997; Thorndyke & Hayes-Roth, 1982). This means that people who learn the layout of very-large-scale VEs from local maps may develop their spatial knowledge in a manner which was previously associated with learning by unaided navigation.

The scale of a local map may be such that the presence of small objects in a VE can be detected, but the exact nature of the objects can not be determined (of course, as is the case with the global map, this may be seen in people's normal, "human's-eye" view of the environment), or the map may use an even larger scale so that objects' shape and identity may also be determined at the expense of displaying even less of the VE. However, one advantage of a local map over a global map is that people may detect objects whenever they are within the map's bounds, even if the people are facing in another direction. Therefore, they are less likely to travel past objects without noticing them.

A navigational aid may also consist of both local and global maps. This type of aid may combine the advantage of the global information offered by the global map with the advantages of local context and object detection offered by the local map, providing the benefits are not outweighed by the attentional demands of the increased amount of navigational information (i.e., two maps rather than one).

Other types of aids include landmarks, grids, numerical coordinates and a compass. Landmarks and grids must be incorporated within the VEs and, therefore, change the environments themselves, something that may not be acceptable in all situations. Numerical coordinates and a compass may be combined to form an aid that does not alter environments, and provides both position and orientation information (the basic information that is required for efficient navigation), although the position information is not pictorially realistic. Moreover, this combination of aids does provide sufficient information to allow people to use lawnmower search strategies.

The following experiments investigated the effects of different navigational aids when participants navigated very-large-scale virtual seascapes. The first experiment used five display conditions: (a) no aids, (b) a global map, (c) a local map, (d) local and global maps (the L&G map) , and (e) numerical coordinates and a compass. All the maps were displayed using a constant, north-up orientation and were dynamic YAH maps which used visual momentum to indicate participants' position, direction of view and direction of travel.

The first condition was chosen as a control condition to investigate how effectively participants could search the VEs without any supplementary aids. The map conditions were chosen to investigate the effects of different scales (see above). The global map's scale showed the whole of a VE but in a small amount of detail, the local map's scale was chosen to simulate a situation where small objects could be detected but not uniquely identified, and the L&G map was a display which combined the information that was available in the global and local map conditions. Darken and Sibert (1996a) did not find a clear advantage of a map over other navigational aids. For this reason the fifth condition was included, and the coordinates and compass were chosen in preference to landmarks and a grid because the former give precise position and orientation and are added to the display rather than the geometric structure of an environment itself.

Experiment 1

In Experiment 1 participants used each of the five display conditions to perform an uninformed search and then an informed search. In the uninformed searches participants had to search previously unknown VEs for nine objects, and in the informed searches participants had to revisit five of the objects they had previously found. The environments and search tasks were based on those used by Darken and Sibert (1993, 1996a), and the metrics used to determine participants' spatial knowledge were similar to those used in earlier studies that investigated navigation in large-scale virtual buildings (Ruddle et al., 1997).

Method

Participants

A total of 10 participants took part in the experiment. They were divided into five groups, which each contained one man and one woman. Each participant underwent a practice session and then navigated five different virtual seascapes using the five different display conditions identified above. The display condition used to navigate each seascape, and the order in which participants used the display conditions and the seascapes were counterbalanced using a Latin Squares design. All the participants were either undergraduates or graduates, who volunteered for the experiment and were paid an honorarium for their participation. Their ages ranged from 19 to 38 years (M = 24.7). Two participants asked to withdraw during the practice session and were replaced in the experiment.

Materials

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

The seascapes were designed to be similar to those used by Darken and Sibert (1993, 1996a). Each of the six seascapes (one practice seascape and five test seascapes) measured 100 X 100 arbitrary units and contained a blue texture mapped “sea” and four “islands”, the outlines of which were copied from an atlas. Each island’s terrain was green and was made up of a flat interior, 250 units above sea level, and cliffs which dropped vertically into the sea.

Each seascape contained three-dimensional models of the same nine everyday objects (a bowl, a cake, a car, a clock, a cup (a mug), a house, a cooking pot (a saucepan), a toaster, and a truck), each of which had a familiarity rating of at least four out of five (Snodgrass & Vanderwart, 1980). The objects’ position was different in each seascape and was determined using random offsets from a 3 X 3 grid. The objects were always situated at sea level and on the sea. No two objects could be seen simultaneously either in the environment or on the local map (see below). All the objects were scaled so their size was approximately 0.07 X 0.07 X 0.07 units. This meant the objects were nearly invisible when seen from a distance of seven units and Performer’s level of detail facility was used to automatically cull objects from the display if they were at a greater range than this (this caused a slight visual “popping” effect if participants looked at an object when it crossed the culling range). The distance of 7 units was determined during a pilot study. When a larger distance (25 units) and correspondingly larger objects were used, participants found the objects quickly by chance, even in the no aids condition.

To define what was seen on the monitor, the application had to specify the height above the sea at which viewing took place. Navigation has been found to be easier when people were allowed to use an elevated ("bird's-eye") view, particularly in situations where no supplementary aids were provided (Darken & Sibert, 1993). Therefore, participants' view height in the present study was elevated and set to 0.4 units above sea level, equivalent to the maximum height allowed by Darken and Sibert (1996a). The center of the screen was marked using a small green square (the target-sight).

An interface which allowed participants to look around while traveling in either a straight line or a curved path, was provided by using the mouse and five keys on the keyboard. The mouse controlled the view direction in three ways: (a) By holding down the left or right mouse buttons, a full 360° horizontal rotation could be performed, (b) By using the mouse to move the cursor up and down the screen, the vertical view direction (pitch) could be changed by ± 90°, and (c) By using the mouse to move the cursor away from the center of the screen the horizontal view direction changed by an amount which increased with the mouse's horizontal offset from the screen center. Method (c) also caused a participant's direction of travel to change and, therefore, allowed them to travel in a curved path. Four of the keys allowed the participant to slow down, stop, speed up, and move at the maximum allowed speed (5 units/s). The fifth changed the participant's direction of movement to the current view direction. All participants mastered this interface without difficulty. Participants were prevented from traveling beyond the edge of the VE by a collision detection algorithm.

Display conditions. Figures 1 and 2 show examples of the screen layouts that were used for the display conditions. The layouts consisted of a large viewport which showed a human's-eye view of the VE and either zero, one, or two smaller viewports which contained maps. The horizontal FOV of the human's-eye viewport was the same for all five conditions (45°), and corresponded to the angle subtended by the monitor when seen from a normal viewing distance (50 cm). The vertical FOV of this viewport was calculated from the horizontal FOV, using the viewport's aspect ratio. A horizontal, red arrow was displayed just above the sea, pointing in a direction that corresponded to participants' direction of movement.

In the no aids condition, the human's-eye viewport covered the whole of the screen (1280 X 1024 pixels; the display shown in Figure 1, but without the coordinates and the compass). In the L&G map condition, the human's-eye viewport covered an area of 1024 X 1024 pixels, and the local and global maps each covered an area of 256 X 256 pixels on the right-hand side of the screen (see Figure 2). This arrangement was chosen in preference to positioning the maps "inside" the VEs (suspended just in front of participants' position, see Darken & Sibert, 1996a) so that the maps did not obscure any part of the VE. The difference between the size of the human's-eye viewport in the no aids and map conditions is characteristic of the compromises that have to be made when overview diagrams are included on computer displays.

The displays used for the local map and global map conditions were the same as the display used for the L&G map condition, except that the global and local map respectively, were not shown. The coordinates/compass condition used the same display layout as the no aids condition, except that the participants' X and Y coordinates were displayed in the top right-hand corner of the screen and a compass was displayed inside the VE, at the bottom of the screen (see Figure 1).

The local and global maps were designed using visual momentum in the form of a moving, green triangle (a wedge; see Aretz, 1991). The global map covered the same area as the whole VE (100 X 100 units), showed the islands, but did not show the positions of the nine objects. The map's orientation (north-up) and position was fixed, which meant that participants appeared to move across the map. One vertex of the triangle indicated participants' position, and the bounds of the triangle indicated where objects would be visible in the human's-eye viewport (i.e., a range of 7 units), given the objects' size and participants' current (momentary) orientation. A small red arrow, one end of which marked participants' position, indicated their direction of movement. In the L&G map condition, a green square indicated the area covered by the local map. The position and orientation of the green triangle, red arrow and green square were all updated at the same rate as the application (20 Hz).

The local map also had a constant, north-up orientation. It was displayed with participants' position at its center, which meant that the environment appeared to move, not the participant (i.e., the map's position was ego-referenced but its orientation was world-referenced). The positions of the nine objects were marked using small, yellow crosses, which were visible when they were within the map's bounds (14 X 14 units). The bounds were chosen so that the crosses became visible on the local map at the same range as the objects themselves became visible in the human's-eye viewport. As in the global map, a green triangle and a red arrow were used to continuously indicate participants' position, horizontal FOV and direction of movement.

Procedures

Participants were run individually and were told that the experiment was being performed to assess people's spatial knowledge in VEs. Each participant underwent a practice session, which lasted approximately 1.5 hr, and then two test sessions, which each lasted approximately 3 hr. The test sessions were always performed during a single week.

Practice procedure. During the practice session, participants were first allowed to become familiar with the controls using the practice VE. Then participants performed a complete practice VE test using the practice seascape, followed by a practice plan-view test. The practice VE test was split into uninformed search and informed search stages. The display used for this practice contained all the aids (i.e., the local map, global map, numerical coordinates and the compass). At the start of this practice, the experimenter placed a sheet of paper in front of participants, which named the nine objects, described what they looked like, specified the width of the seascape (100 units) and the field of view of the human's-eye viewport (45°). Participants initial position was beside one of the objects and they searched the VE for all nine objects in any order. A list of the objects to be found was displayed on the screen and when participants found an object they centered the target-sight over the object and pressed the "y" key. This caused the object to be removed from the list, but the object itself remained displayed within the VE. When participants had found the ninth object they started the informed search stage.

During the informed search participants revisited five of the objects in a specific order, using the knowledge they had gained about the VE's layout, and at each object made estimates of direction (compass bearing; the VE-orientation data) and of straight-line distance (VE-straight data) to the other eight objects in the VE. At all times, the name of the next object to be found was displayed on the screen. Participants indicated they had found the object by centering the target-sight over the object and pressing "y" key. This caused the VE software to move participants until they were directly above the object. The participant then estimated the direction to each of the other objects by rotating their viewpoint until they thought they were facing the object and indicating this by pressing the "y", key which caused the view direction to be recorded. Then a Motif (UNIX) window was presented eight times and participants entered estimates for the straight–line distance to each of the other objects. Once all the direction and distance estimates had been performed a message was displayed on the screen which named the next object to be revisited. Thus the order of objects to be revisited was only revealed one object at a time.

The practice plan-view test was a computerized version of the paper plan-view test used in studies that investigated spatial knowledge development in virtual and real buildings (Ruddle et al., 1997; Thorndyke & Hayes-Roth, 1982) and took approximately 10 min to complete. For the computerized version, 40 similar screens were displayed. Figure 3 shows that the position of two of the objects was marked on each screen and the name of a third object was displayed in the top left-hand corner. The initial orientation of the objects was random but, by holding down the left or right mouse buttons, the participant could rotate the display to any other orientation. The scale of all the displays was the same. For each screen participants moved the mouse until they thought the cursor lay over the position of the third object, relative to the other two objects, and pressed the space bar. This caused the objects estimated position to be recorded and the next screen to be displayed.

Test procedure. During the two test sessions participants performed five tests, each using a different seascape and a different display condition. Each test comprised a VE test followed by a plan-view test, and these followed the same procedure as the practice VE test and the practice plan-view test described above. The tests for two of the display conditions were performed in one session and the tests for the remaining three display conditions were performed during the other session. Participants' were given unlimited time for each VE test and their movements through the VEs were recorded continuously for later analysis. However, an hour after the start of each test the experimenter offered to intervene and take participants to the remaining objects.

Results

Preliminary Data Analysis

Only the data from the five test display conditions was analyzed. The efficiency of participants' uninformed searches was measured by calculating the total distance they traveled and the time they took to locate the nine objects. The minimum distance participants had to travel to conduct an exhaustive search of one of the seascapes and, therefore, be sure of finding all nine objects, was approximately 700 units. If they always moved at the maximum possible speed, this took 140 s. All except one participant successfully completed the uninformed search in each test. This participant had still not found one of the objects after browsing for an hour during the coordinates/compass condition. At this point the experimenter intervened and "took them" to the object. To analyze the data this participant's uninformed search time and distance for this condition were set equal to the maximum of any participant when using any of the aids. Participant means for the distance and time measures correlated very highly, $r = .88$, $p < .01$, so the distance traveled was used for the statistical analyses.

The efficiency of participants' informed searches was measured by computing the distance they traveled to visit all nine objects in excess of the minimum possible distance as a percentage of the minimum, the percentage extra distance traveled (PE-distance). A similar calculation was computed using time instead of distance, with the minimum possible time calculated by dividing the minimum possible distance by the maximum speed of movement (5 units/s). Participant means for the two measures correlated very highly, $r = .97$, $p < .01$, so the PE-distance metric was used for the statistical analyses. Each participant's appreciation of the relative distances between the objects was calculated by correlating their VE-straight distance estimates with the corresponding distances between objects in the VE, and normalizing the correlations' using Fisher's r–to–z transformation. Each participant's direction estimate accuracy was determined by calculating the mean angular error of their VE-orientation estimates.

All the participants successfully completed the informed search when using the L&G map and the local map, but the experimenter intervened with one participant in the global map condition, and five participants each in the coordinates/compass and no aid conditions. In situations where this occurred participants' PE-distance data were set equal to the maximum values of any participant when using any of the aids.

As in the studies by Thorndyke and Hayes-Roth (1982) and Ruddle et al. (1997), each participant's plan view test data was used to calculate the mean angular error of their estimates, relative to the two marked locations on each of the 40 screens, and the correlation of the distances between the participant's estimated object positions and the actual distances. The distribution of the latter data was then normalized using Fisher's r-to-z transformation.

Uninformed Search

The distribution of the distances participants traveled during their uninformed searches was normalized using a logarithmic transformation. A repeated measures analysis of variance (ANOVA), illustrated in Figure 4, showed there was a main effect of display condition on the distances participants traveled, $F(4, 36) = 7.46$, $p < .01$.

Planned contrasts, summarized in Table 1, showed that participants traveled significantly shorter distances when using either the L&G map or the local map than when using the global map, the coordinates/compass, or no aid. However, even when using the L&G map, participants traveled slightly more than 70% farther than the minimum possible distance.

Informed Search

The distribution of participants' PE-distance data was normalized using a logarithmic transformation. A repeated measures ANOVA, illustrated in Figure 5, showed there was a main effect of display condition on PE-distance, $F(4, 36) = 12.85$, $p < .01$. Planned contrasts, summarized in Table 2, showed that participants traveled significantly less extra distance when using the L&G map than in any other of the conditions, and traveled significantly shorter distances when using either the local map or the global map than when using the coordinates/compass or no aid. Even when using the L&G map, participants traveled an average of 80% farther than the minimum possible distance.

The distribution of participants' mean VE-orientation estimate errors was normalized using a logarithmic transformation. A repeated measures ANOVA, illustrated in Figure 6, showed there was a main effect of condition on the accuracy of participants' estimates, $F(4, 36) = 22.81$, $p < .01$. Planned contrasts, summarized in Table 3, showed that participants made more accurate estimates when using either the L&G map or the global map than when using the local map, the coordinates/compass, or no aid.

Participants' sense of the relative VE-straight distances between the objects was analyzed using their Fisher's z data. A repeated measures ANOVA, illustrated in Figure 7, showed a similar pattern of results to the VE-orientation data. There was a main effect of display condition on participants' sense of relative distance, $F(4, 36) = 12.58$, $p < .01$, and planned contrasts, summarized in Table 4, showed that participants sense of relative distance was more accurate when they used either the L&G map or the global map than when they used the local map, the coordinates/compass, or no aid

Plan View Test

This test was designed to compare the effects of the different display conditions on participants' ability to transform their spatial knowledge to the perspective used to display the maps (a plan-view). A repeated measures ANOVA, illustrated in Table 5, showed there was a main effect of display condition on participants mean direction estimate error, $F(4, 36) = 4.44$, $p < .01$. Planned contrasts showed that participants' direction estimates were significantly more accurate when they used the global map than when they used the coordinates/compass, $F(1, 9) = 13.40$, $p < .01$, or no aid, $F(1, 9) = 10.91$, $p < .01$, and were significantly more accurate when they used the local map than when they used the coordinates/compass, $F(1, 9) = 4.98$, $p < .05$. None of the other differences was significant.

A repeated measures ANOVA, illustrated in Table 6, showed there was a main effect of display condition on participants sense of relative distance, $F(4, 36) = 3.38$, $p < .05$. Planned contrasts showed that participants' sense of relative distance was significantly more accurate when they used the global map than when they used the local map, $F(1, 9) = 5.37$, $p < .05$, the coordinates/compass, $F(1, 9) = 12.66$, $p < .01$, or no aid, $F(1, 9) = 6.07$, $p < .05$. None of the other differences was significant.

Search Strategies

As well as analyzing navigation performance and spatial knowledge, we investigated the strategies that participants used to search for the objects. To do this we wrote a second Performer application that displayed the paths traveled by each participant, overlaid on a plan view of the appropriate VE. The shape of the paths varied widely and this meant that our attempts to determine the search strategy associated with each path were highly subjective. However, the search strategy data do provide an insight into how, rather than how effectively, participants searched and, therefore, are a useful supplement to the main body of data described above.

Initial analysis of the paths indicated that participants regularly used only four identifiable search strategies. These are illustrated in Figure 8 and were: (a) edge, (b) lawnmower, (c) via object, and (d) direct. In edge searches participants followed the edges of the islands and the VE boundaries, and lawnmower searches were characterized by movement in straight, parallel and perpendicular lines. Via object and direct searches were only applicable to the informed searches, for which participants already had knowledge of the target objects' location. Via object searches were where participants travelled to an object(s) they had previously visited and then travelled directly to the target object, and direct searches were those in which participants travelled in an approximately straight line directly to the target objects. Participants' paths which did not fit any of the above four search strategies were classified as other. Once the above common strategies had been determined, two judges independently rated the paths by awarding one point to the strategy that most closely matched each participant's path to each object during the uninformed and informed searches. Despite the subjective nature of the rating process, the judges agreed on the strategies associated with 78% of the paths although a large number of these agreements were classified as other.

Table 7 shows that participants frequently used an edge strategy for their uninformed searches but rarely used a lawnmower strategy, despite the fact that this type of search represented a very efficient way of locating all the objects. Table 8 shows that participants' strategies varied more between the different display conditions for the informed search than for the uninformed search, and this may reflect differences in the spatial knowledge that participants had developed. With the L&G map participants frequently remembered the position of objects and moved directly to them. The post-processing software was also used to determine the number of times participants missed objects (traveled past an object they were searching for without detecting it; see Table 9). The most common situation in which this occurred was when participants were not provided with the local map and were facing away from the object when they traveled past.

Discussion

Overall, and as expected, the L&G map was the most effective display condition, although the local map was equally effective when used for the uninformed search. When performing the informed search with the L&G map, participants traveled approximately 1/3 of the distance that they traveled when using the local map or the global map, and approximately 1/11 of the distance that they traveled when using the coordinates/compass or no aid. However, even with the L&G map participants traveled 80% farther than necessary, although their direction and distance judgments indicate that they knew the approximate locations of the objects.

In keeping with the findings of Darken & Sibert (1996a), edge following was the most commonly used strategy in the uninformed search and participants seemed to perform this by anchoring their searches around recognizable features in the VEs (the island and VE boundaries). Following a lawnmower pattern would have been much more efficient but was rarely used. Given the difficulty that participants had finding the objects in conditions such as the coordinates/compass, it was surprising that they did not use a lawnmower strategy more often. In fact, one participant commented that they tried to find one particular object using this strategy but was unable to do so! This suggests that people may need to be trained to perform lawnmower searches.

The plan view test investigated participants' ability to transform their knowledge of the objects' position to a plan view (map) perspective. Participants' direction and relative distance estimates in this test were expected to be most accurate with the global map and the L&G map because in both of these conditions participants were directly shown their global position each time they found an object. In fact, there were fewer significant differences between conditions in the plan view test than for the corresponding data in the VE tests. This may be attributed to additional sources of experimental error which each added to variance in the data, namely the facts that: (a)

The test measured participants' knowledge after they had left a VE, (b) The test screens used a different scale to that the global map, (c) The screens did not show the islands' outline (reference information which was shown on the global map), and (d) Tiredness caused by the length of each test session (approximately 3 hr).

One intriguing result was that, although participants traveled shorter distances in the L&G map condition than in the global map condition, their distance and direction estimates were slightly more accurate in the latter for both the VE and the plan view tests. A possible explanation is that participants learned the position of the objects more accurately in the global map condition by necessity, because they had such difficulty finding them (see Table 9).

By contrast, although participants traveled similar distances in the local and global map conditions, the inferior VE-orientation and VE-straight performance for the local map show that in this condition participants had only a vague idea of objects' relative positions. This echos the situation that is found when people learn spatial knowledge by unaided navigation (see Ruddle et al., 1997).

Another unexpected outcome was that no significant differences occurred between the coordinates/compass and no aid conditions. In fact, most of the data indicated that participants performed marginally worse when using the coordinates/compass. In both of these conditions participants experienced difficulty finding the objects in the informed search but it should be noted that the mean values for their VE-orientation and VE-straight data were more accurate than those that would have been produced by random guesses (respectively, 90° and zero).

Experiment 2

Experiment 1 demonstrated the overall superiority of the L&G map when participants search unfamiliar, very-large-scale VEs but left unanswered questions that concerned the local and global maps. In particular, would extended navigation with the global map on its own ultimately prove superior to the L&G map, as suggested by the direction and distance estimate data, and would participants' knowledge that was acquired using a local map continue to develop in a manner that is usually associated with learning from unaided navigation?

For Experiment 2 the three map conditions were used. In each condition, participants searched a VE for nine objects (an uninformed search), and then searched the same VE three times for all the objects (informed searches). The coordinates/compass condition was not used because Experiment 1 showed that it was not an effective navigational aid.

Method

Participants

A total of 12 participants (four male and eight female) took part in the experiment and were randomly assigned to one of three groups. Each participant underwent a practice session and then navigated three different virtual seascapes using the three display conditions identified above. The display condition used to navigate each seascape, and the order in which participants used the display conditions and the seascapes were counterbalanced using a Latin Squares design. All the participants were either undergraduates or graduates, were different to the participants who took part in Experiment 1, volunteered for the experiment and were paid an honorarium for their participation. Their ages ranged from 18 to 39 years (M = 21.3).

Materials

The experiment was performed using the same hardware, software application and interface as Experiment 1. The practice VE was the same as that used in Experiment 1 and the test VEs were three of those used in Experiment 1. The display conditions used the same screen layouts as the corresponding conditions in Experiment 1.

Procedures

Participants were run individually and were told that the experiment was being performed to assess people's spatial knowledge in VEs. Each participant underwent a practice session, which lasted up to 2 hr, and then three test sessions, which each lasted up to 3 hr. The test sessions were always performed during a single week.

Practice procedure. During the practice session, participants were first allowed to become familiar with the controls using the practice VE. Then participants performed a two-stage practice test, using a display which contained both types of map that were used in the three test sessions (i.e., the L&G map, see Figure 2).

The first stage of the practice was the same as Experiment 1's uninformed search and when participants had found the ninth object the VE software stopped. The second stage of the practice was similar to the informed search stage of Experiment 1. Participants searched for all nine objects in a specific order, and indicated they had found each object by centering the target-sight over the object and pressing "y" key. Participants' initial position was set so that they were beside the first of the nine objects that had to be found. At the first, third, fifth, seventh and ninth (last) object, participants performed direction and distance estimates for the other eight objects, using the same procedure as in Experiment 1 (the VE-orientation and VE-straight data).

Test procedure. Participants used a different display condition and a different VE in each of the three test sessions. Each session was split into four stages. In the first two stages, participants searched a VE for the nine objects in any order. In the third stage, participants searched the VE for the nine objects in a particular order and performed direction and distance estimates at five of the objects. The fourth stage was the same as the third stage, except that the objects were searched for in a different order. Participants' initial position was different for each stage, but was always beside one of the objects (in stages three and four, this was always the first object that was to be searched for). Participants were allowed up to 45 min to complete each stage. If they had not completed a stage by this time the experimenter intervened and took them to the remaining objects, where they performed the direction and distance estimates (Stages 3 and 4 only). Then the participants continued with the remaining stages of the test.

Results

Any-Order Searches (Test Stages 1 and 2)

One participant had not found one of the objects after 45 min during Stage 1 of their global map search. To analyze the data this participant's search time and distance for this test stage were set equal to the maximum of any participant when using any of the display conditions for Stages 1 or 2. Participant means for the distance and time measures correlated very highly, $r = .87$, $p < .01$, so the distance traveled was used for the statistical analyses.

The distribution of the distances participants traveled during their searches was normalized using a logarithmic transformation. As expected a repeated measures ANOVA, illustrated in Figure 9, showed there were main effects of display condition and stage number on the distances participants traveled (respectively, $F(2, 22) = 35.15$, $p < .01$, and $F(1, 11) = 21.33$, $p < .01$). Planned contrasts, showed that in both stages participants traveled significantly shorter distances in the L&G map condition than in the local map condition (Stage 1, $F(1, 11) = 4.47$, $p < .05$; Stage 2, $F(1, 11) = 6.32$, $p < .05$), or the global map condition (Stage 1, $F(1, 11) = 27.84$, $p < .01$; Stage 2, $F(1, 11) = 14.86$, $p < .01$). Participants traveled significantly shorter distances in the local map condition than in the global map condition in Stage 1, $F(1, 11) = 10.00$, $p < .01$, but this difference was not significant for Stage 2.

Specific-Order Searches (Test Stages 3 and 4)

Three participants in Stage 3 of the local map, two participants each in Stage 4 of the local map and Stage 3 of the global map, and one participant in Stage 4 of the global map, had not found all the objects after 45 min. To analyze the data these participants' search distance time and distance for these test stages were set equal to the maximum of any participant when using any of the display conditions for Stages 3 and 4. Participant means for the distance and time measures correlated very highly, $r = .85$, $p < .01$, so the distance traveled was used for the statistical analyses.

The percentage extra distance (PE-distance) that participants traveled was calculated in the same way as in Experiment 1, and the distribution of these data was normalized using a logarithmic transformation. As expected a repeated measures ANOVA, illustrated in Figure 10, showed there were main effects of display condition and stage number on participants' PE-distance data (respectively, $F(2, 22) = 20.53$, $p < .01$, and $F(1, 11) = 16.48$, $p < .01$). Planned contrasts showed that participants traveled significantly shorter distances in the L&G map condition than in the local map condition for Stage 3, $F(1, 11) = 52.24$, $p < .01$, and Stage 4, $F(1, 11) = 34.32$, $p < .01$, or in the global map condition for Stage 3, $F(1, 11) = 7.83$, $p < .05$, but this difference was not significant for Stage 4. Participants also traveled significantly shorter distances in the global map condition than in the local map condition for Stage 3, $F(1, 11) = 19.62$, $p < .01$, and Stage 4, $F(1, 11) = 22.10$, $p < .01$.

The distribution of participants' mean VE-orientation data was normalized using a logarithmic transformation and analyzed using a repeated measures ANOVA. Figure 11 shows that there were main effects of display condition and stage number on the accuracy of participants' VE-orientation judgments (respectively, $F(2, 22) = 44.01$, $p < .01$, and $F(1, 11) = 11.31$, $p < .01$). Planned contrasts showed that participants made significantly more accurate VE-orientation judgments in the L&G and global map conditions than in the local map condition in Stage 3 (respectively, $F(1, 11) = 91.54$, $p < .01$, and $F(1, 11) = 148.67$, $p < .01$) and in Stage 4 (respectively, $F(1, 11) = 85.90$, $p < .01$, and $F(1, 11) = 103.91$, $p < .01$). Participants made more accurate VE-orientation judgments in the global map condition than in the L&G map condition in Stage 3, $F(1, 11) = 6.89$, $p < .05$, but this difference was not significant for Stage 4.

Participants' sense of the relative distances between the objects was analyzed using their Fisher's z data. A repeated measures ANOVA, illustrated in Figure 12, indicated a similar pattern of results to participants' VE-orientation data. There were main effects of display condition and stage number on the accuracy of participants' sense of relative distance (respectively, $F(2, 22) = 48.23$, $p < .01$, and $F(1, 11) = 16.37$, $p <$

.01). Planned contrasts showed that participants had a significantly more accurate sense of relative distance in both the L&G and global map conditions than in the local map condition in Stage 3 (respectively, $F(1, 11) = 83.32$, $p < .01$, and $F(1, 11) = 111.78$, $p < .01$) and in Stage 4 (respectively, $F(1, 11) = 98.18$, $p < .01$, and $F(1, 11) = 90.53$, $p < .01$). The differences between the L&G and global map conditions were not significant for either stage.

Search Strategies

Two judges used the software and points system used in Experiment 1 to independently rate the strategies that participants used to find the objects in each test stage. The judges agreed on the strategies used for 79% of the searches and, as in Experiment 1, these data are included to supplement the distance traveled, distance estimate and direction estimate data by helping to illustrate how participants searched.

Three participants used the lawnmower strategy for nearly all of their searches in Stage 1 and this was the cause of the increased popularity of landmark searches when compared with the uninformed searches that were carried out in Experiment 1 (see Tables 7 and 10). The frequency of use of the strategies changed from Sessions 2 to 4 because participants knowledge of the object's positions improved. In Session 4 of the L&G map and global map conditions, participants traveled directly to approximately 40% of the objects but in the local map condition participants also frequently used lawnmower and via object strategies. Table 11 shows that, as in Experiment 1, participants missed objects much more frequently in the global condition, than in either the L&G map, or the local map condition.

Discussion

As in Experiment 1, the L&G map was the most effective navigational aid, although the global map held a slight, but significant, advantage for the VE-orientation estimates that participants carried out in Stage 3, an advantage which was in line with the distance and direction estimate data of Experiment 1. In Stage 2 of the L&G map condition, participants traveled a shorter distance than was necessary to search the whole environment and this indicated that they had already developed a reasonable idea of where the objects were located.

The greatest improvement across the test stages occurred in the global map condition. In Stage 1, participants travelled almost three times farther than in the L&G map condition, but in Stage 4 there was no significant difference between the distances participants traveled in these two conditions, or the accuracy of participants' direction and distance estimates. Although participants still frequently missed objects in their searches during Stage 4 of the global map condition, they often seemed to realize this quickly, returned to objects' approximate position, searched again and located them without traveling a large, extra distance.

Participants improved less quickly in the local map condition. Although they were able to quickly find the objects in the uninformed search (Stage 1), they had much more difficulty than in the other conditions finding the objects in the specific-order searches. One reason for this difference may be that participants learned objects' absolute (world-referenced) positions in the global map and the L&G map conditions, whereas in the local map condition participants often learned the position of objects in relation to each other.

General Discussion

Overall and as has already been stated, the L&G map was the most effective navigational aid. It allowed participants to learn the positions of objects in a world-referenced frame so they could be easily relocated, and also displayed participants' immediate surroundings on a map that allowed objects to be detected once participants were in the objects' approximate vicinity.

However, after sufficient time participants were able to find objects as efficiently in the global map condition as in the L&G map condition (see Experiment 2, Stage 4). Contrary to expectations, the direction and distance estimate data from both experiments consistently show that participants had slightly more accurate knowledge of where objects were in the former condition than in the latter. This difference is important but the reason is open to speculation. As has already been suggested, the most likely reason is that participants learned the position of the objects more accurately in the global map condition by necessity, and some supporting evidence comes from the percentage of times that participants missed the objects (see Tables 9 and 11). Another possibility is that the provision of two maps in the L&G map condition increased the attentional demands made on participants and this caused a decrease in the accuracy of their estimates.

The global map presented information in a fixed, world-referenced position and orientation. The local map used a fixed, world-referenced orientation but an ego-referenced position (at any time it showed the immediate surroundings of a participant). This meant that participants' had to integrate information that was learned from many different positions to develop their spatial knowledge, and this may explain why their knowledge developed more slowly than in the other map conditions. The effect of map scale on participants knowledge in the present study has similarities with the differences in spatial knowledge that occur when people either perform unaided navigation of a large-scale environment, or learn the environment's layout from a single view of a map (see Ruddle et al., 1997; Thorndyke & Hayes-Roth, 1982). In addition, if the local map had used a forward-up orientation then participants are likely to have developed their spatial knowledge even more slowly because they would have had to integrate information that was learned from many different positions <u>and</u> orientations. This may help to explain the differences found by Aretz (1991), Barfield et al. (1995), and Wickens et al. (1996) when participants performed world-centered tasks using either a forward-up, or a north-up map.

One unexpected outcome in Experiment 1 was that participants were marginally worse in coordinates/compass condition than in the no aids condition. Participants were expected to have difficulty locating the objects in the latter, but to use the basic position and orientation information provided by the coordinates and compass to perform systematic searches of the environments in the former. This suggests that coordinates and a compass, or other similar aids such as a virtual sun, are not a suitable for the navigation of very-large-scale VEs unless people are trained to effectively use them.

In situations where there is only sufficient space for one map the data from the present study suggest that the optimum scale of the map depends on people's knowledge of the VE. If people have little (or no) knowledge of the VE a local map is likely to be a more effective aid than a global map. However, once people develop some knowledge of the VE the information provided by a global map allows people to travel through the environment and search more efficiently than when using a local map. On the other hand, designers of environments for computer games and puzzles that involve spatial exploration may wish to only provide a local map to deliberately restrict the rate at which users develop knowledge of VEs' layout.

In keeping with the findings of Darken and Sibert (1996a) most participants used an edge strategy more often than any other strategy to perform the uninformed searches. Even in the controlled setting of the laboratory, where there was a specific task to perform, they tended to use this strategy in preference to more efficient strategies such as a lawnmower strategy. It seems that participants preferred to anchor their searches around the major topological features in the environment (the islands and environment edges), rather than search throughout the environment. This has implications for VE designers, who may consider placing the locations they want people to visit near major topological features, or displaying certain features on maps of VEs to influence where people are likely to travel. Within applications such as virtual shopping malls, these manipulations could be used to promote prime locations for virtual shops, in the same way as shops situated in the central parts of towns typically attract a higher rent than shops situated in less frequented places.

Finally, the VEs used in the present study allowed participants unconstrained movement and, therefore, they could move directly to any position. In some VE applications people's movement is constrained to routes that are defined by solid barriers such as walls. In these VEs, detailed route information is likely to have increased importance and can be clearly displayed on a local map.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

References

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

Barfield, W., Rosenberg, C., & Furness, T. A. III (1995). Situation awareness as a function of frame of reference, computer-graphics eyepoint elevation, and geometric field of view. International Journal of Aviation Psychology, 5, 233-256.

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 (CHI'96), 142-149.

Ellis, S. R., Smith, S., Grunwald, A., & McGreevy, M. W. (1991). Direction judgement error in computer generated displays and actual scenes. In S. R. Ellis, M. K. Kaiser, & A. J. Grunwald (Eds.), Pictorial communication in virtual and real environments (pp. 504-526). London: Taylor & Francis.

Encarnação, J., Foley, J., Bryson, S., Feiner, S. K., & Gershon, N. (1994). Research issues in perception and user interfaces. IEEE Computer Graphics and Applications, 14, 67-69.

Engineering and Construction. (1997, March). VR News, 6, 26-30.

Headmounted displays. (1997, May). VR News, pp 28-33.

Hooper, E. Y., & Coury, B. G.. (1996). Graphical displays for orientation information. Human Factors, 36, 62–78.

Johnson, S. L., & Roscoe, S. N. (1972). What moves, the airplane or the world? Human Factors, 14, 107–129.

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

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

Lynch, K. (1960). The image of the city. Cambridge: MIT Press.

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.

Research Directions in Virtual Environments. (1992). Computer Graphics, 26, 153-177.

Roscoe, S. N. (1968). Airborne displays for flight navigation. Human Factors, 10, 321–332.

Roscoe, S. N., Corl, L., & Jensen, R. S. (1981). Flight display dynamics revisited. Human Factors, 23, 341-353.

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. (1997). Navigating buildings in "desk-top" virtual environments: Experimental investigations using extended navigational experience. Journal of Experimental Psychology: Applied.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1998). Navigating large-scale "desk-top" virtual buildings: Effects of orientation aids and familiarity. Presence: Teleoperators and Virtual Environments, 7, 179-192.

Siegel, A. W. (1981). The externalization of cognitive maps by children and adults: In search of ways to ask better questions. In L. Liben, A. Patterson, & N. Newcombe (Eds.), Spatial representation and behavior across the life span (pp. 167-194). 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.

Spatially immersive display systems. (1998, May). VR News, pp 23-29.

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.

Virtual Tourism: Paris. (1996). Cupertino, CA: Apple Computer. [Computer Software].

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., Liang, C. -C., Prevett, T., & Olmos, O. (1996). Electronic maps for terminal area navigation: Effects of frame of reference and dimensionality. International Journal of Aviation Psychology, 6, 241-271.

Wickens, C. D., & Prevett, T. T. (1995). Exploring the dimensions of egocentricity in aircraft navigation displays. Journal of Experimental Psychology: Applied, 1, 110–135.

Williams, Hutchinson, & Wickens, C. D. (1996) A comparison of methods for promoting geographic knowledge in simulated aircraft navigation. Human Factors, 38, 50-64.

Woods, D. D. (1984). Visual momentum: A concept to improve the cognitive coupling of person and computer. International Journal of Man-Machine Studies, 21, 229–244.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

Table 1

Planned Contrasts for Distance Traveled in Uninformed Searches. * p < .05; ** p < .01.

Display condition	L&G map	Local map	Global map	Coordinates/compass
Local map, F(1, 9)	0.80	--	--	--
Global map, F(1, 9)	11.01**	5.87*	--	--
Coordinates/compass, F(1, 9)	21.19**	13.74**	1.65	--
No aid, F(1, 9)	12.28**	6.80*	0.03	1.21

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

Table 2

Planned Contrasts for Percentage Extra Distance (PE-distance) Traveled in Informed Searches. * p < .05; ** p < .01.

Display condition	L&G map	Local map	Global map	Coordinates/compass
Local map, F(1, 9)	12.85**	--	--	--
Global map, F(1, 9)	9.75**	0.21	--	--
Coordinates/compass, F(1, 9)	36.34**	5.97*	8.44**	--
No aid, F(1, 9)	38.50**	6.87*	9.50**	0.03

Table 3

Planned Contrasts for VE-orientation Estimate Errors. * p < .05; ** p < .01.

Display condition	L&G map	Local map	Global map	Coordinates/compass
Local map, F(1, 9)	20.35**	--	--	--
Global map, F(1, 9)	1.45	32.65**	--	--
Coordinates/compass, F(1, 9)	38.71**	2.93	55.13**	--
No aid, F(1, 9)	30.36**	1.00	45.06**	0.51

Table 4

Planned Contrasts for VE-straight Distance Correlations, Transformed from Fisher's z. * p < .05; ** p < .01.

Display condition	L&G map	Local map	Global map	Coordinates/compass
Local map, F(1, 9)	16.36**	--	--	--
Global map, F(1, 9)	0.12	19.27**	--	--
Coordinates/compass, F(1, 9)	21.93**	0.41	25.28**	--
No aid, F(1, 9)	19.52**	0.14	22.70**	0.07

Table 5

Participants' Plan View Test Mean Direction Estimate Errors

Display condition	L&G map	Local map	Global map	Coordinates /compass	No aid
Mean (degrees)	31.3	30.5	25.9	37.6	36.5
SE	3.4	3.1	2.9	2.3	1.5

Table 6

Participants' Plan View Test Mean Distance Correlations,

Display condition	L&G map	Local map	Global map	Coordinates /compass	No aid
Mean z	0.58	0.51	0.85	0.33	0.49
SE	0.16	0.16	0.10	0.09	0.06
Equivalent r	0.53	0.47	0.69	0.32	0.46

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

Table 7

Percentage of Uninformed Searches for which each Search Strategy was used

Strategy	L&G map	Local map	Global map	Coordinates /compass	No aid
Edge	38	46	34	34	40
Lawnmower	1	1	5	4	1
Other	37	35	47	41	38
Help given	0	0	0	1	0
Judges disagree	24	18	14	20	21

Table 8

Percentage of Informed Searches for which each Search Strategy was used

Strategy	L&G map	Local map	Global map	Coordinates /compass	No aid
Edge	0	16	4	6	10
Lawnmower	0	4	0	0	0
Via object(s)	6	6	0	4	2
Direct	56	12	25	8	10
Other	14	33	45	18	27
Help given	0	0	2	33	29
Judges disagree	24	29	24	31	22

Table 9

Percentage of times Objects were Missed during Uninformed and Informed Searches

% missed	L&G map	Local map	Global map	Coordinates /compass	No aid
Uninformed search	2	2	44	36	21
Informed search	2	2	27	15	22

Table 10

Percentage of Searches for which each Strategy was used in Experiment 2. Dashes indicate that a Strategy was not applicable to that Stage.

Condition/	L&G map				Local map				Global map			
Stage	1	2	3	4	1	2	3	4	1	2	3	4
Edge	18	15	4	1	17	10	4	2	14	9	1	1
Lawnmower	35	5	2	2	26	12	12	11	30	2	0	0
Via object(s)	--	0	3	3	--	1	10	13	--	2	2	1
Direct	--	41	45	43	--	32	7	21	--	24	38	41
Other	30	26	22	31	37	26	14	20	36	42	37	37
Help given	0	0	0	0	0	0	16	5	1	0	6	1
Judges disagree	17	13	24	20	20	19	37	28	19	21	16	19

Table 11

Percentage of times Objects were Missed during Searches in Experiment 2

Condition	L&G map	Local map	Global map
Session 1	0	5	42
Session 2	1	3	39
Session 3	2	2	23
Session 4	1	0	21

Figure Captions

Figure 1.

Screen display used in the compass/coordinate display condition.

Figure 2.

Screen display used in the L&G map display condition.

Figure 3.

Screen display used in the plan view test.

Figure 4.

Distance traveled during the uninformed search of each display condition in Experiment 1. Crds/comp = coordinates/compass. Error Bars indicate MSE.

Figure 5.

Percentage extra distance traveled during the informed search of each display condition in Experiment 1. Crds/comp = coordinates/compass. Error Bars indicate MSE.

Figure 6.

Participants' mean VE-orientation estimate errors for each display condition in Experiment 1. Crds/comp = coordinates/compass. Error Bars indicate MSE.

Figure 7.

Participants' mean distance correlations for each display condition in Experiment 1, transformed from their mean z values. Crds/comp = coordinates/compass. Error Bars indicate MSE.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied, 5*, 54-75.

Figure 8.

Examples of the search strategies used in Experiment 1. Clockwise, from top left, they are edge, lawnmower, via object and direct.

Figure 9.

Distance traveled during the any-order searches of each display condition in Experiment 2. Error Bars indicate MSE.

Figure 10.

Percentage extra distance traveled during the specific-order searches of each display condition in Experiment 2. Error Bars indicate MSE.

Figure 11.

Participants' mean VE-orientation estimate errors for Stages 3 and 4 of each display condition in Experiment 2. Error Bars indicate MSE.

Figure 12.

Participants' mean distance correlations for Stages 3 and 4 of each display condition in Experiment 2, transformed from their mean z values. Error Bars indicate MSE.