

T. Todd Elvins

todd@acm.org

Indicast Corporation

Carlsbad, CA 92008

David R. Nadeau

nadeau@sdsc.edu

San Diego Supercomputer Center

Rina Schul

rschul@crl.ucsd.edu

SDSU/UCSD Joint Doctoral

Program in Clinical Psychology

David Kirsh

kirsh@cogsci.ucsd.edu

Cognitive Science Department

University of California, San Diego

Worldlets: 3-D Thumbnails for Wayfinding in Large Virtual Worlds

Abstract

Finding one's way to sites of interest on the Web can be problematic, and this difficulty has been recently exacerbated by widespread development of 3-D Web content and virtual-world browser technology using the Virtual Reality Modeling Language (VRML). Whereas travelers can often navigate 2-D Web sites based on textual and 2-D thumbnail image representations of the sites' content, finding one's way to destinations in 3-D environments is notoriously troublesome. Wayfinding literature provides clear support for the importance of landmarks in building a cognitive map and then using that map to navigate in a 3-D environment, be it real or virtual. Textual and 2-D image landmark representations, however, lack the depth and context needed for travelers to reliably recognize 3-D landmarks. This paper describes a novel 3-D thumbnail landmark affordance called a *worldlet*. Containing a 3-D fragment of a virtual world, worldlets offer travelers first-person, multi-view-point experience with faithful representations of potential destinations.

To facilitate an investigation into the comparative advantages of landmark affordances for wayfinding, worldlet capture algorithms were designed, implemented, and incorporated into two VRML-based virtual environment browsers. Findings from a psychological experiment using one of these browsers revealed that, compared to textual and image guidebook usage, worldlet guidebook usage: nearly doubled the time subjects spent studying the landmarks in the guidebook, significantly reduced the time required for subjects to reach landmarks, and reduced backtracking to almost zero. These results support the hypothesis that worldlets facilitate traveler landmark knowledge, expedite wayfinding in large virtual environments, and enable skilled wayfinding.

I Introduction

As the Web moves into 3-D, the problem of finding one's way to sites of interest is exacerbated. URLs currently are represented on a page or menu by sentences as in figure 1, or by 2-D thumbnail images that depict some aspect of the URL as in figure 2. It is reasonable to suppose that interactive 3-D thumbnails—or *worldlets*—will soon come into use. To test whether worldlets improve wayfinding, we implemented a prototype 3-D system and designed a psychological experiment aimed at answering two questions:

- Are menus of worldlets better than menus of text and 2-D images for wayfinding?

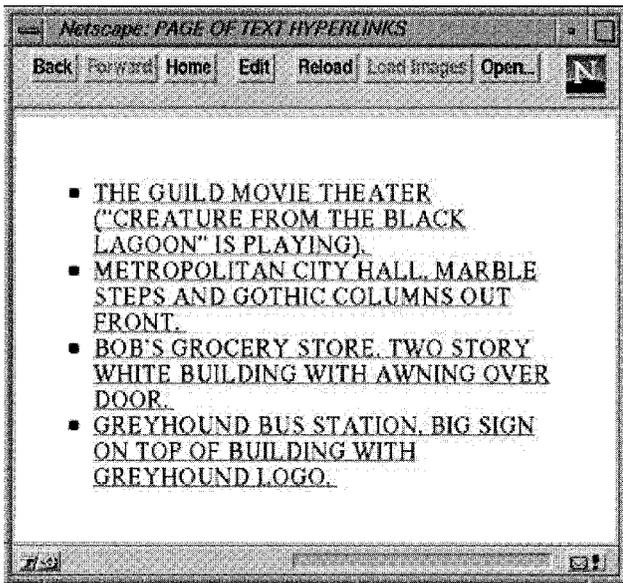


Figure 1. A Web page linking to alternative destinations within a virtual world. On this page, the destinations are represented with text.

- What aspects of the wayfinding process are most enhanced by the use of worldlets?

In the experiment, subjects were placed at starting locations in three virtual cities and asked to find their way to goal locations via a series of landmarks. In one city, the landmarks were described with text, in another city the landmarks were described using 2-D thumbnail images, and in yet another city the landmarks were described using manipulable worldlets, as shown in figure 3. We hypothesized that, compared to text and thumbnail images, worldlets would facilitate better landmark knowledge and expedite wayfinding.

2 Wayfinding

Wayfinding is “a term that can refer to a rather narrow concern: That is, how well people are able to find their way to their particular destination without delay or undue anxiety” (Peponis, Zimring, & Choi, 1990, p. 561). Whereas *navigating* is most often defined as the process of following a course by making directed movements, wayfinding is the critical process of determining

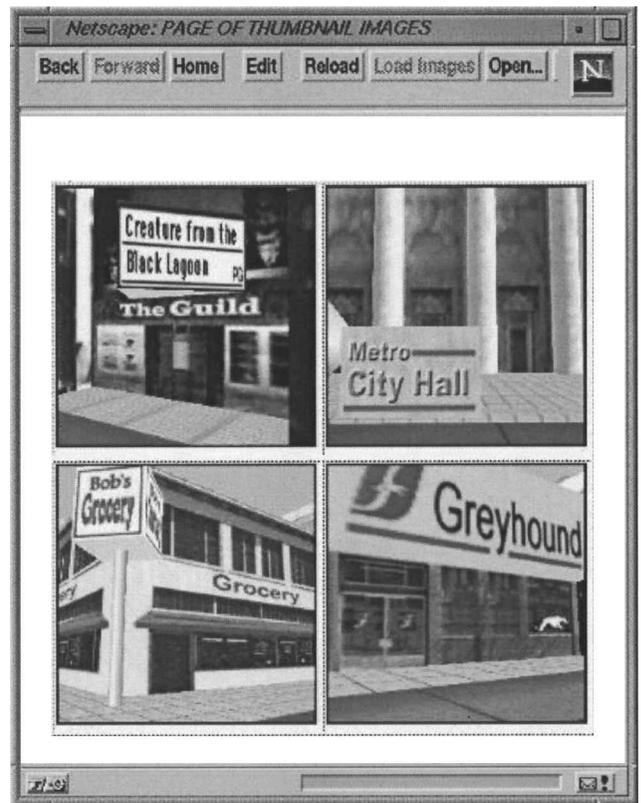


Figure 2. On this Web page, the destinations from Figure 1 are represented as 2-D thumbnail images.

the strategy, direction, and course needed to reach a desired destination. Travelers find their way using survey, procedural, and landmark knowledge (Goldin & Thorndyke, 1982; Peponis et al., 1990). Each type of knowledge helps the traveler construct a cognitive map of a region and thereafter find their way using that map (Appleyard, 1976; Downs & Stea, 1973; Lynch, 1960; Passini, 1992).

Survey knowledge provides a map-like, bird's-eye view of a region and contains spatial information including locations, orientations, and sizes of regional features. Procedural knowledge characterizes a region by memorized sequences of actions that construct routes to desired destinations. Landmark knowledge records the visual features of landmarks, including their 3-D shape, size, texture, and so forth (Appleyard, 1976; Goldin &



Figure 3. Construction of a worldlet. A gas station landmark in a virtual city (top left) viewed from a viewpoint in the street (top right). Viewed from above, the gas station is upper left, and a region around the previous viewpoint is highlighted (bottom left). The gas station is captured within a spherical worldlet, and (bottom right) is viewed from street level within the worldlet.

Thorndyke, 1982). For a structure to be a landmark, it must have high imagability; that is, it must be distinctive and memorable (Lynch, 1960).

Landmarks are the subject of landmark knowledge, but they also play a part in survey and procedural knowledge. In survey knowledge, landmarks provide regional anchors with which to calibrate distances and directions. In procedural knowledge, landmarks mark decision points along a route, helping in the recall of procedures to get to and from destinations of interest.

Whereas a traveler's landmark knowledge characterizes a destination by its 3-D shape, size, texture, and so forth, traditional virtual environment (VE) browser menus and landmark guidebooks characterize destinations only by textual descriptions or images. This representation mismatch reduces the effectiveness of destination menus and guidebooks. Unable to engage their memory of 3-D landmarks to recognize destinations of interest, travelers may resort to a naive, exhaustive search to find a desired landmark.

2.1 Related Work

Although social and engineering scientists have studied human wayfinding in the real world for many years (Appleyard, 1976; Downs & Stea, 1973; Goldin & Thorndyke, 1982), only recently have computer scientists investigated this subject. Satalich (1995) led subjects on a tour of a virtual building and then tested the subjects' ability to find their way to places seen during the tour. Darken and Sibert (1996) applied Lynch's (1960) city structure theories to large, sparsely featured VEs and then compared the subjects' ability to find their way to target destinations with and without the added city structure. Witmer, Bailey, Knerr, and Parsons (1996) studied route learning and found that subjects who rehearsed a route in a virtual building learned the route better than did subjects who rehearsed the route verbally. Gillner and Mallot (1997) found that subjects pieced together survey knowledge when repeatedly required to find shortest paths between a dozen unfamiliar locations in a VE. In studying VEs for real-world task training, Waller, Hunt, and Knapp (1998) found that long exposure fostered good spatial knowledge under several VE fidelity conditions. Ruddle, Payne, and Jones (1998) studied path complexity and a compass condition as factors in VE route learning and found that, beyond two turns, subjects were quickly disoriented with or without a compass. Although each of these studies recognized the value of landmark knowledge for wayfinding, none studied the value of 3-D landmarks in familiarizing subjects with an environment before entering it. This is a primary goal of our study.

2.2 Wayfinding Affordances

Three-dimensional models of large cities are becoming available for exploration via the Web (Art+Com, 1997; Liggett, Friedman, and Jepson, 1995; University College London, 1997). As discussed previously, travelers visiting such large VEs for the first time are easily disoriented, may have difficulty identifying a place upon arrival, and may not be able to find their way back to a place just visited (Downs & Stea, 1973; Darken & Sibert, 1996). The following subsec-

tions describe affordances that travelers currently use to scrutinize potential landmark and VE destinations.

- *Textual descriptions* are the dominant method used to represent VE landmarks in world- and viewpoint-selection user interfaces. HTML pages, for instance, often provide lists of available Web-based worlds (such as those authored in the Virtual Reality Modeling Language (VRML) (ISO, 1997)), each one characterized by a URL, a VE name, and/or a brief description. Within VRML worlds, textual descriptions characterize viewpoints and describe destinations associated with clickable anchor shapes.
- *Icons, thumbnail images, and image maps* provide common visual wayfinding aids. In a 3-D context, games often provide "jump gates" onto which images of remote destinations are texture mapped. Stepping through such a gate jumps the traveler to the destination landmark depicted on the gate.
- *Image mosaics* group together multiple captured images into a traversable structure. Apple's QuickTime VR, for instance, can use images captured from multiple viewing angles at the same viewing position (Chen, 1995). By ordering images within a traveler-centered cylindrical structure, QuickTime VR can provide a traveler the ability to look in any direction through automatic selection of an appropriate image from the structure. By chaining multiple mosaic structures together, the content author can create a walkthrough path that hops from viewpoint to viewpoint. Similar image mosaics can be used to create zoom paths, pan paths, and so forth.
- *Miniature Worlds and Maps* are available in most 3-D environment browsers and computer games, and they enable the traveler to zoom out and view the world in miniature and to thereby gain survey knowledge. Pausch, Burnette, Brockway, and Weiblen (1995) extend this notion by creating a world in miniature (WIM) embedded within the main world. The miniature world duplicates all elements of the main world and adds an icon denoting the traveler's position and orientation. Held within the traveler's virtual hand, the traveler can reach into the miniature and reposition world content

and/or themselves. Simultaneously, the outer main world is updated to match the altered miniature, automatically adjusting the positions of shapes or the traveler. Similarly, 2-D and 3-D maps are frequently found as navigation aids within VEs. For instance, 3-D games often provide a 3-D reduced-detail map in which an icon denotes the player's location. Such maps can be panned, zoomed, and rotated to provide alternate viewpoints similar to that which is possible with miniature worlds.

None of these wayfinding affordances offer the traveler faithful 3-D landmark representations or a first-person experience with potential destinations. Wickens and Baker (1995) argue that VE search tasks are difficult because "often the object of the search is specified in a format different from and more abstract than its rendering in the VE". We address this problem by representing the object of the search in the same 3-D format as the VE.

3 What are Worldlets?

In previous papers, we described our ongoing worldlet wayfinding research (Elvins, 1997; Elvins, Nadeau, & Kirsh, 1997; Elvins, 1998; Elvins, Nadeau, Kirsh, & Schul, 1998). Similar to a 3-D photograph, a worldlet is associated with a viewing position and orientation, and captures a snapshot of the 3-D shapes (usually landmarks) that fall within the viewpoint's viewing volume. Worldlets can be explored and manipulated in the same way that virtual worlds can be explored and manipulated. Incorporated into a VRML browser, a guidebook or list of worldlets enables a world traveler to interactively view landmarks and landmark context from multiple viewpoints. Seeing the landmarks presented in worldlets helps a traveler select destinations of interest and find their way in unfamiliar worlds.

In typical use, the worldlet's viewpoint is aimed at an important landmark, and the worldlet's captured shapes reconstruct that landmark and its associated context. When viewed within an interactive 3-D browser, a



Figure 4. The navigation browser's viewpoint selection window and the wayfinding browser's worldlet guidebook. The "Go to" button was removed in the wayfinding browser.

worldlet provides a manipulable 3-D thumbnail representation of the landmark.

Clicking on a worldlet in a landmark menu might jump the traveler to the landmark represented by the worldlet. The landmark could be at a viewpoint in the current world or in another world. A menu of worldlets, such as that shown in figure 4, might resemble landmark history lists, bookmarks, or travel guidebooks. Travelers are free to study the landmark in the worldlet from any viewpoint before deciding whether to go there.

We have developed two types of worldlets:

- A *frustum worldlet* contains shapes within a standard pie-shaped viewing frustum, positioned and oriented based upon a selected viewpoint. When viewed, a frustum worldlet looks like a pie-shaped fragment clipped from the larger main world.
- A *spherical worldlet* contains shapes within a spherical viewing bubble, positioned at a selected viewpoint with a 360 deg. field of view. When displayed, a spherical worldlet looks like a ball-shaped world fragment, similar to a snowglobe knick knock depicting a favorite attraction.

For both worldlet types, hither and yon clipping planes restrict the extent of the worldlet, ensuring that the worldlet contains a manageable subset of the larger main world. Worldlet shape content is shaded and textured to match the corresponding shapes in the main world. Although the main world may have content that changes over time, the captured worldlet remains static, recording the content of the world at the time the worldlet was captured.

Figure 3 depicts the process of constructing a spherical worldlet from a fragment of a virtual city surrounding a central viewpoint position. Once captured, worldlets can be browsed, manipulated, and explored in a number of ways.

4 Worldlets in Browsers

We have incorporated worldlets into the user interface for two experimental VRML browsers: a navigation browser, and a wayfinding browser. The navigation browser provides a “Go to” feature that enables travelers to study, select, and fly to world viewpoints in the current world, and to study, select, and load previously visited worlds on the browser’s history list. Guidebooks of worldlets in the wayfinding browser allow travelers to study 3-D landmark guidebooks before and during travel in an unfamiliar VE. No go-to button is present in this browser in which travelers manually move themselves about based on the landmark knowledge gleaned during study.

4.1 Navigation Browser

Traditional VRML browsers provide a viewpoint menu that offers a choice of viewpoints, each denoted by a brief textual description. We have extended this standard feature to provide three experimental viewpoint selection interfaces, each using worldlets. All three present a set of worldlets, one for each author-selected viewpoint in the world. A “Go to” button flies the main window’s viewpoint to that of the currently selected worldlet. The browser also supports on-the-fly capture of worldlets from the traveler’s current viewpoint.

- The *viewpoint list window* provides a list of worldlets beside a worldlet viewer. Selection of a worldlet from the list displays the worldlet in the viewer where it may be interactively panned, zoomed, and rotated.
- The *viewpoint guidebook window* presents a grid of worldlet viewers, arranged to form a guidebook photo-album page. Buttons on the window advance the guidebook forward or back a page at a time. Each worldlet on the page can be interactively examined. Figure 4 shows the viewpoint guidebook window.
- The *viewpoint overlay window* enables the traveler to select a worldlet from a list, and overlay it atop the main window, highlighted in green. This worldlet overlay provides a clear indication of the worldlet’s viewpoint position and orientation, along with the portion of the world captured within that worldlet. Figure 3b was generated using this overlay technique.

Traditional VRML browsers provide a history list of recently visited worlds, each denoted by its title or URL. We have extended this standard feature to provide two world selection interfaces, each using worldlets. In each world selection window, a “Go to” button loads into the main window the world associated with the currently selected worldlet.

- The *world list window* provides a list of worldlets beside an interactive worldlet viewer, similar to the viewpoint list window discussed earlier. One worldlet is available for each world on the browser’s history list.
- The *world guidebook window* uses the same guidebook photo-album layout used for the viewpoint guidebook window discussed earlier. One worldlet is available for each world on the history list.

A “Save as” feature of the VRML browser enables the traveler to save a worldlet to a VRML file. Using a collection of saved worldlets, a world author can create a VRML world of worldlets. Such a world acts like a 3-D destination index, similar to a shelf full of snowglobe knick knacks. When cast as a VRML anchor shape, a

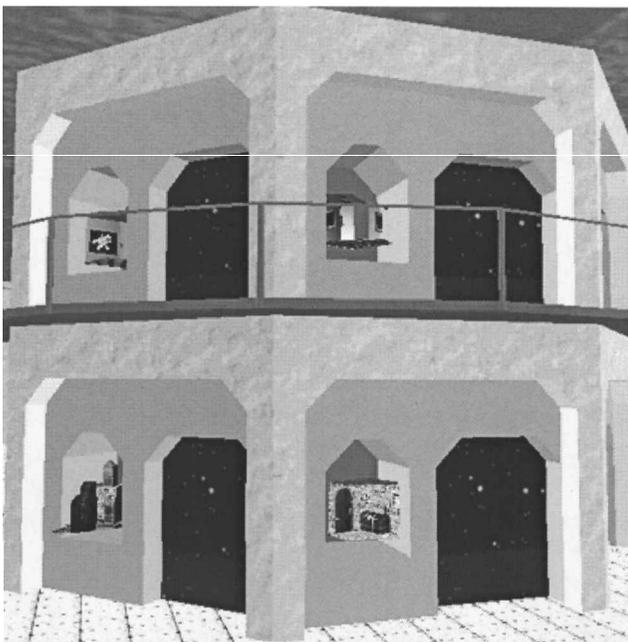
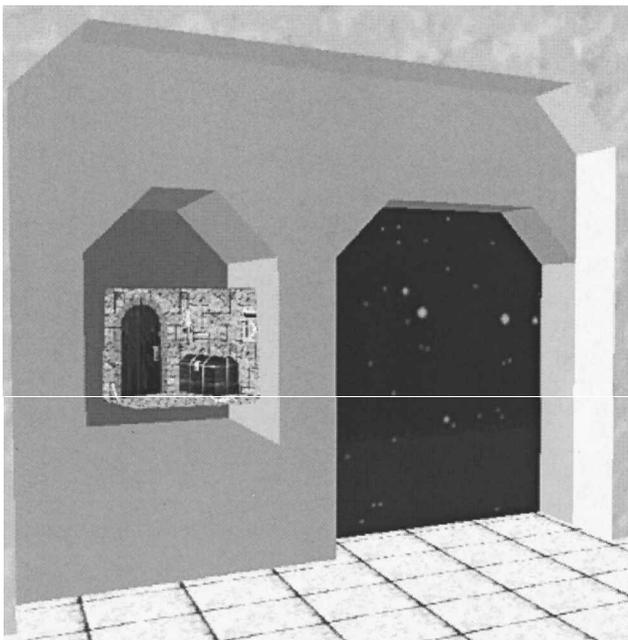


Figure 5. A world of worldlets that (top) associates a worldlet with each doorway (bottom) in an environment containing multiple such doorways. Each doorway leads to a different world.

worldlet provides a 3-D button that, when clicked on, loads the associated world into the traveler's browser.

Figure 5 shows a world of worldlets. Figure 5a shows

a close-up view of a world "doorway" and a niche that contains a worldlet illustrating a viewpoint in that world. Figure 5b shows a wider view of the same world and multiple such doorways.

4.2 Wayfinding Browser

The second browser was implemented specifically for the psychological experiment. After manipulating worldlets in a landmark guidebook to familiarize themselves with the 3-D size, shape, color, texture, and context of the landmarks, subjects in the experiment found their way manually in the main window based upon their landmark knowledge. The landmark guidebook resembled the viewpoint selection window shown in figure 4, but it did not offer a "Go to" button. The wayfinding browser is described in detail in the methods section.

5 Implementation

Worldlets can be implemented in several ways. The implementation presented here is based on a hardware-accelerated algorithm and support algorithms developed specifically for fast execution. The VRML browser used in this work maintains VE geometry within a tree-like scene graph. Worldlets are also stored as scene graphs, together with additional state information. To capture a worldlet or to display a worldlet or VE, the VRML browser traverses the associated scene graph and feeds a 3-D graphics pipeline.

5.1 Worldlet Capture in General

Any 3-D graphics pipeline can be roughly divided into two stages: (i) transform, clip, and cull, and (ii) rasterize (Foley, van Dam, Feiner, & Hughes, 1990). The first stage applies modeling, viewing, perspective, and viewport transforms to map 3-D shapes to the 2-D viewport. Along the way, shapes outside of the viewing frustum are clipped away and backfaces are removed. The second stage uses 2-D shapes output by the first

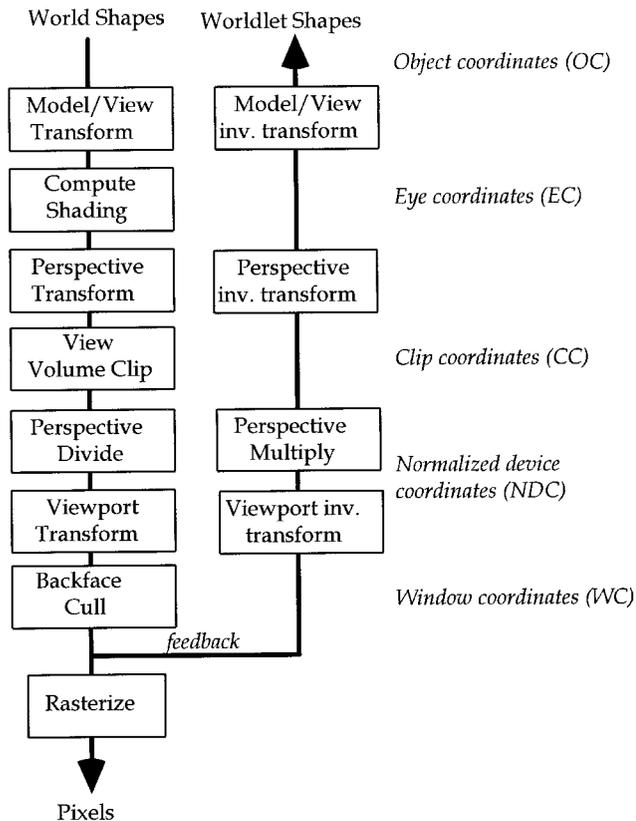


Figure 6. World shapes are transformed to screen space via the downward OpenGL graphics pipeline on the left. Worldlet shapes are inverse transformed from window coordinates to object coordinates via the upward pipeline in the middle of the figure. The intermediary coordinate systems are named on the right.

stage and draws the associated points, lines, and polygons on the screen.

Worldlet capture taps into this 3-D graphics pipeline, extracting the transformed, shaded, clipped, and culled shape coordinates output by the first stage prior to rasterization in the second stage. An extracted coordinate contains x and y screen-space components, a depth-buffer z -space component, and the w coordinate. Each extracted coordinate has an associated RGB color and texture coordinates, which are computed by shading and texture calculation phases in the first pipeline stage.

To create a worldlet, these extracted coordinates are untransformed to map them from viewport space back to world space as shown in figure 6. The inverses of the

viewport, perspective, viewing, and modeling transforms are each applied. Coordinate RGB colors and texture coordinates are used to reconstruct 3-D worldlet geometry in a worldlet scene graph.

Display of a worldlet passes this 3-D geometry back down the graphics pipeline, transforming, clipping, culling, and rasterizing the worldlet like any other 3-D content.

5.2 Frustum and Spherical Worldlets

A frustum worldlet is the result of capturing 3-D graphics pipeline output for a single traversal of the scene graph as viewed from the traveler's current viewpoint. The shape set extracted after the first pipeline stage contains only those points, lines, and polygons that fall within the viewing frustum. The worldlet constructed by the browser from this geometry looks like a pie-shaped slice cut out of the world.

A spherical worldlet is the result of performing multiple frustum captures and combining the results. The VRML browser captures a spherical worldlet by spinning about a viewpoint position, generating a set of frustum worldlets that each use a different viewing orientation. Additional captures aimed straight up and straight down complete the spherical worldlet. The resulting set of capture geometry constructs a 360 deg. spherical view from the current viewpoint.

When displayed, the spherical worldlet's geometry looks like a bubble cut out of the VE. A close yon clip plane keeps the bubble small, ensuring that it captures only landmark features in the immediate neighborhood and not the entire virtual world.

5.3 Worldlet Capture in OpenGL

To take advantage of the rendering speed offered by the accelerated 3-D graphics pipeline on high-speed workstations, we implemented worldlet display and capture using OpenInventor (Wernecke, 1994) and OpenGL graphics libraries from Silicon Graphics. Scene graph construction and display traversal is managed by OpenInventor. To capture worldlet geometry, the VRML browser places the pipeline into feedback mode

prior to a capture traversal, and returns it to rendering mode following traversal.

While in feedback mode, the OpenGL pipeline diverts all transformed, clipped, and culled coordinates into a buffer provided by the browser. Upon completion of a capture traversal, no rasterization has taken place, and the feedback buffer contains the extracted geometry. By parsing the feedback buffer, the VRML browser reconstructs worldlet geometry, applying appropriate inverse transforms.

OpenGL feedback buffer information includes shape coordinates, colors, and texture coordinates, but it does not include an indication of which texture image to use for which bit of geometry. To capture this additional information, the VRML browser uses OpenGL's pass-through features to pass custom flags down through the pipeline during traversal. To prepare these pass-through flags, the browser augments the world scene graph prior to traversal, assigning a unique identifier to each texture image. During a capture traversal, each time a texture image is encountered, the associated identifier is passed down through the pipeline and into the feedback buffer along with shape coordinates, colors, and texture coordinates. During parsing of the feedback buffer, these texture identifiers enable worldlet geometry reconstruction to apply the correct texture images to the correct shapes.

6 Methods

The psychological experiment was designed to test whether guidebooks of worldlets are better for wayfinding than guidebooks of text and images, and, if so, why.

6.1 General Description

A standard, within-subject, randomized design was used. Training began by reading each subject a series of instructions and teaching the subject, in a practice virtual city (three blocks by three blocks), how to operate the user interface. Subjects were allowed to spend unlimited time in this city becoming familiar with the controls and the online guidebooks.

Each subject was then asked to find his or her way, via one landmark, to the goal (a brightly colored kiosk) in three practice cities, each five blocks by five blocks. In one city, the guidebook represented each landmark with a paragraph of text, in another city the landmarks were represented as 2-D thumbnail images, and in a third city the landmarks were represented as spherical worldlets. The guidebooks contained only landmark information, and subjects were given no procedural or survey information. Although the image and worldlet guidebook had the same appearance, as shown in figure 4, the worldlets in the worldlet guidebook could be explored and manipulated.

After completing the practice tasks, the experiment began, and subjects were asked to find their way to the goal kiosk via eight ordered landmarks in three virtual cities, ten blocks by ten blocks. Performance in these three cities was timed, and the subject's location, orientation, and actions were recorded at 1 sec. intervals. Subjects visited the test cities in a random order to counterbalance possible carryover effects, and each of the cities and landmarks were different and unique.

Before the computer portion of the experiment, subjects completed a brief questionnaire on their prior experience with computers. After the experiment, subjects answered written questions about the strategies they used while finding their way, and how difficult they found the task using each of the guidebooks. Subjects also completed spatial and verbal neuropsychological tests during a separate session.

6.2 Constants

The major independent variable in this experiment was the type of guidebook used (text, image, or worldlet). Parameters concerned with city design and landmark layout, as listed in table 1, were held constant.

6.3 City Design

Our primary goal in designing the virtual cities was to preserve as many elements of a real city as possible. Toward this goal, we developed a pseudo-random city generator that produced cities based on a regular

Table 1. *Although the Arrangement of Buildings, Landmarks, and Routes Were Different in Each City, Their Basic Structure in Terms of City Design and Landmark Layout was Kept Constant to Allow for Controlled Comparison. Subjects Moved at a Fixed Speed in all Cities*

Parameter	Constant value	Comments
City dimensions	ten by ten blocks	100 total blocks
Building lots per block	twenty on perimeter of block	building footprint varies
Distance from start to goal	1,200 meters	past 100 12-meter-wide lots
Landmarks per city	nine, including goal kiosk	ordered in guidebooks
Distance between landmarks	1.5 blocks	± 0.25 block
Travel speed	five meters/second	slow driving speed
Types of landmarks	eight	one of each kind in each city
Turns required to reach goal	five	left and right turns
Radius of information	approximately twenty meters	see city design section
Landmark positions	three on corner, five midblock	goal is midblock
Goal building	yellow/red striped kiosk	six-sided, viewpoint invariant

street grid with pavement roads and sidewalks between the blocks. Each block contained up to twenty buildings, side by side around the block's perimeter. Using a cache of 250 building designs, buildings were randomly selected and placed on city blocks. Many non-landmark buildings were repeated between and within the cities. Rendering requirements were reduced by programming buildings to be invisible when farther than three blocks from the current viewpoint. To prevent subjects from seeing buildings appearing and disappearing in the distance, we added fog to the environment.

Cities contained parks, parking lots, and other civic features but did not contain cars or street signs. Landmarks were placed so that a subject could not short-cut the route and would never have to make an exhaustive search; as long as they carefully looked down each possible path for familiar landmarks and landmark context, they could always determine in which direction to proceed. No landmarks were located within one block of the perimeter wall that surrounded each city.

6.4 Radius of Information

The radius of information was the most difficult parameter to make constant across the three guide-

books. Given that text, 2-D images, and worldlets are different representations, they do not convey an equivalent quantity and quality of landmark information. We determined, while building the guidebooks for the pilot study, that using a radius of 20 m produced the most nearly equivalent set of guidebooks. For the text guidebook, this meant describing in words the landmark, the building or structure directly across the street, the building on each side, and whether the landmark was on a corner. Thus, our textual description of the landmark shown in figure 3 read:

Citgo Gas Station. Citgo's red gas pumps are at a corner location. Next door is the Cafe 360 in a beige colored building, and across the street is the Lumbini Restaurant in a red brick building.

For the image guidebook, a 20 m radius meant positioning the camera across the street from the landmark so that the captured thumbnail image would contain the buildings or structures on the left and right of the landmark. For the worldlet guidebook, setting the information radius simply meant setting the yon clip plane to cutaway 3-D shapes farther away than 20 m while capturing a spherical worldlet. Image and worldlet camera



Figure 7. *The main city window.*

locations were selected so as not to suggest a direction of approach to the landmark.

6.5 Hardware and Software Design

The experiment was conducted using a standard 19 in. monitor. Stereo glasses were not used. A main city window displayed the city in the wayfinding VRML browser. Subjects used keyboard arrow keys to move forward and backward a fixed distance on each key press, and to turn left and right by a fixed angle on each key press. In the worldlet guidebook, subjects used the arrow keys in the same manner to move around in the worldlet. Collision detection prevented subjects from passing through buildings and other objects in the city.

Subjects pushed a Start button before beginning, a Landmark button each time a landmark was reached, and a Stop button upon reaching the goal kiosk. Subjects pushed a Guidebook button to display the guidebook. Subjects practiced operating all of these buttons in practice cities prior to beginning the experiment. The windows were designed so that, when the guidebook was open (shown in figure 4), the main city window (shown in figure 7), was completely occluded, and no movement could be made.

6.6 Lessons Learned from the Pilot Study

Our pilot study design and results are described by Elvins et al. (1997). From the pilot study, we learned that landmarks cannot be more than two blocks apart because subjects frequently resort to an exhaustive search of the neighborhood. We also learned that some subjects will take much longer than other subjects to complete the task and that the task difficulty and length must be designed while keeping the slower wayfinders in mind.

6.7 Subject Instructions

The following is a subset of the instructions read to the subjects. Instructions 3 and 4 were repeated before beginning each city.

1. The guidebook contains a series of landmarks that you probably want to follow, since traveling past these landmarks will lead you to the goal kiosk.
2. When you are at a landmark, you may have to walk as far as one or two blocks before you will be able to see the next landmark.

3. Your task is to locate the goal kiosk as quickly as possible. Go as fast as you can, but be careful to not get lost.
4. Remember, you can look at the guidebook as often as you like.

6.8 Subjects

Subjects were primarily University of California–San Diego students from a number of major departments. Ages ranged from 18 to 36, with a mean age of 24. Fourteen men and twelve women completed the experiment. Results from four other subjects had to be excluded from our data because the subjects became lost in one or more of the cities. All subjects were computer literate, but with varying degrees of experience with VEs.

6.9 Neuropsychological Tests

To provide an independent measure of a subset of the subjects' cognitive abilities, subjects were tested in a separate session, using two standardized subtests from the Wechsler Adult Intelligent Scale-Revised (WAIS-R) (Wechsler, 1955). The Vocabulary subtest, comprising part of the WAIS-R Verbal subtests, was used to assess the subject's verbal ability. Performance was scored according to standardized procedures based on the subject's degree of familiarity with a list of up to 35 words and their elegance of expression. The second test was the Block Design subtest comprising part of the WAIS-R Performance subtests, which was used to assess the subjects' constructional-spatial ability, and was administered in the standardized manner.

6.10 Analysis

A one-way repeated measures analysis of variance (ANOVA) was performed for each of the dependent variables and the overall times and distances. The within-subjects variable was the landmark description type with three levels: text, image, and worldlet. Post hoc analyses were completed using the Scheffé test for differences in sample means. Correlation analysis was per-

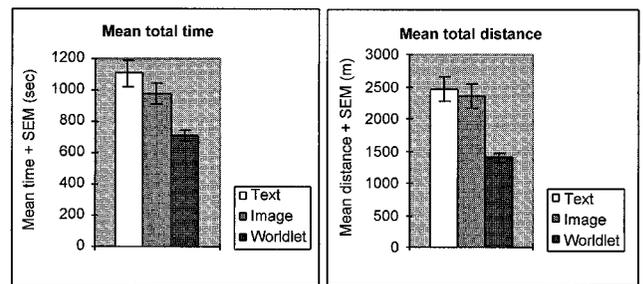


Figure 8. (left) Mean total time. The overall time spent traveling to the goal was significantly less using the worldlet guidebook than when using the text ($F(1, 24) = 9.89$; $p < 0.01$) or image ($F(1, 24) = 4.63$; $p < 0.01$) guidebooks. (right) Mean total distance. The overall distance traveled using the worldlet guidebook was significantly less than the distance traveled using the text ($F(1, 24) = 11.41$; $p < 0.01$) or image ($F(1, 24) = 8.93$; $p < 0.01$) guidebooks.

formed across the scores of the neuropsychological tests and the scores of the wayfinding tests. We adopted a significance level of 0.05 on all analyses unless otherwise noted.

7 Results

Our primary result was that people spend more time consulting the worldlet guidebook than the other guidebooks, but that they require less time to reach the goal when using the worldlet guidebook. The following specific findings tell us more about the component activities that benefitted from worldlet use.

7.1 Time and Distance Measures

Figures 8 through 11 show time and distance measures in which the independent variable (type of guidebook) is displayed as white, gray, and black bars along with the standard error of the mean (SEM). Figure 8 shows the mean overall time and distance required to reach the goal kiosk, and figure 9 breaks the overall time into its components. Figure 10 gives the forward and backtracking distance components. Figure 11 shows the mean number of times that each type of guidebook

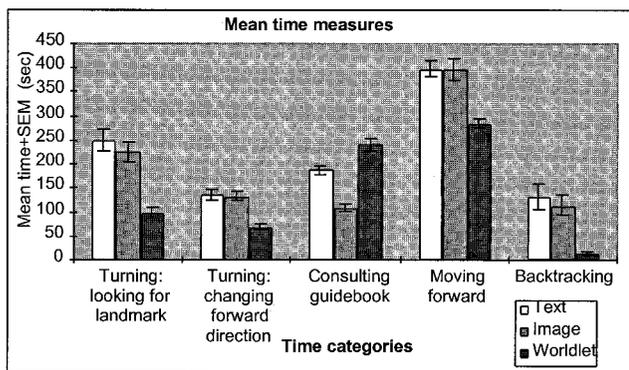


Figure 9. Turning: looking for landmark is time spent standing in one location. Turning: changing forward direction is time spent changing forward direction. Consulting guidebook is time spent with the guidebook on screen. Moving forward is time spent traveling over new ground, and Backtracking is time spent retreating. Note that the time spent consulting the worldlet guidebook is significantly greater than the time spent consulting the text or image guidebook. Also note that the time spent backtracking is near zero when using the worldlet guidebook.

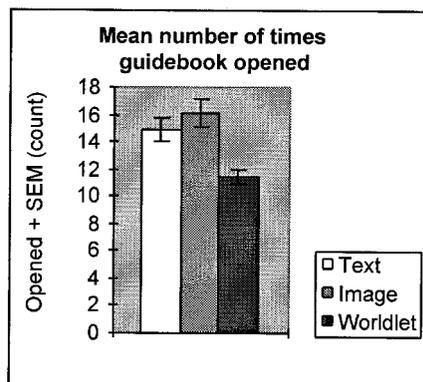


Figure 11. Subjects opened the text and image guidebooks an average of fifteen to sixteen times while seeking the eight landmarks enroute to the goal kiosk. The worldlet guidebook was opened eleven times on average.

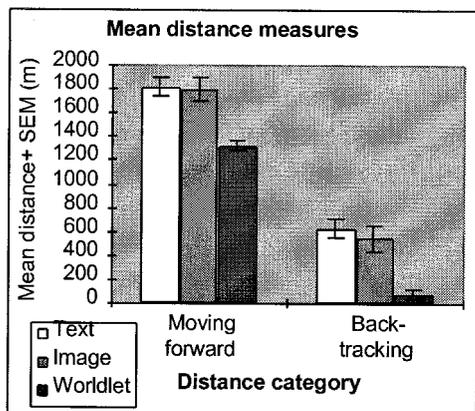


Figure 10. The mean distance moving forward while using the worldlet guidebook is significantly less than the forward distance covered while using the text or image guidebook. Backtracking distance is almost nonexistent for the worldlet guidebook task.

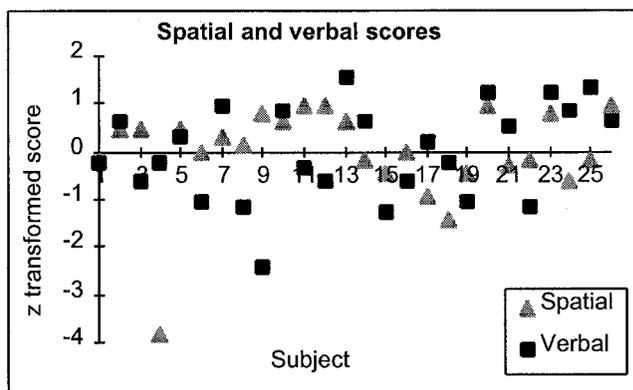


Figure 12. Z-transformed spatial and verbal scores from the neuropsychological tests. The mean raw scores were 45 and 55 for the spatial and verbal tests, respectively.

was opened. Figure 12 diagrams the subjects' scores on the Wechsler neuropsychological tests.

Analysis of variance followed by Scheffé post hoc tests were performed for the different measures. Table 2 summarizes these results.

7.2 Street Plots

Figure 13 shows plots of the streets traversed by the 26 subjects in the three test cities. The thickness and brightness of each street line indicate the total number of times that the street was traversed. For example, wider brighter streets indicate heavily used avenues. Backtracking can increase the number of street traversals beyond the expected minimum of 26. The starting point and goal kiosk are indicated on each plot, and landmarks enroute to the goal are marked with a star (☆).

Table 2. Our ANOVA Indicated Significant Differences between All the Sample Means and Post Hoc Analysis Was Warranted

	ANOVA	Post hoc comparisons (Scheffé <i>F</i> -test)		
		text vs. image	text vs. worldlet	image vs. worldlet
Time				
Consulting guidebook	55.59 * $p < 0.0001$	19.66 * $p < 0.001$	8.84 * $p < 0.01$	54.88 * $p < 0.011$
Looking for landmark	21.61 * $p < 0.0001$	0.47	18.69 * $p < 0.001$	13.25 * $p < 0.01$
Changing heading	17.36 * $p < 0.0001$	0.02	13.53 * $p < 0.01$	12.49 * $p < 0.01$
Moving forward	16.02 * $p < 0.0001$	0.003	12.20 * $p < 0.01$	11.83 * $p < 0.01$
Backtracking	9.55 * $p < 0.0003$	0.22	8.26 * $p < 0.01$	5.86 * $p < 0.01$
Overall	10.34 * $p < 0.0002$	0.98	9.89 * $p < 0.01$	4.63 * $p < 0.01$
Distance				
Moving forward	15.33 * $p < 0.0001$	0.05	12.23 * $p < 0.01$	10.72 * $p < 0.01$
Backtracking	10.19 * $p < 0.0002$	0.23	8.84 * $p < 0.01$	6.21 * $p < 0.01$
Overall	13.65 * $p < 0.0001$	0.15	11.41 * $p < 0.01$	8.93 * $p < 0.01$
Guidebook				
Opened	15.27 * $p < 0.0001$	1.04	7.59 * $p < 0.01$	14.26 * $p < 0.001$

Post hoc (Scheffé) analysis indicated, for all dependent variables, significant differences in the sample means of the text versus worldlet performance, and image versus worldlet performance. Performance using the text versus image guidebooks was not significantly different, except that the time spent consulting the text guidebook was significantly greater than the time spent consulting the image guidebook. In each table cell, the first number is the *F* value, and the second number is the *p* value. Asterisk indicates significance $p < 0.05$. Time and distance measures are explained in the captions for figures 8 through 11.

7.3 Subjective Measures

As part of the written questionnaire administered after the experiment, subjects were asked how difficult the search task was using each of the different types of guidebooks. Although responses were scattered, a clear trend in favor of the worldlet guidebook was evident, as shown in table 3.

7.4 Correlation Analysis

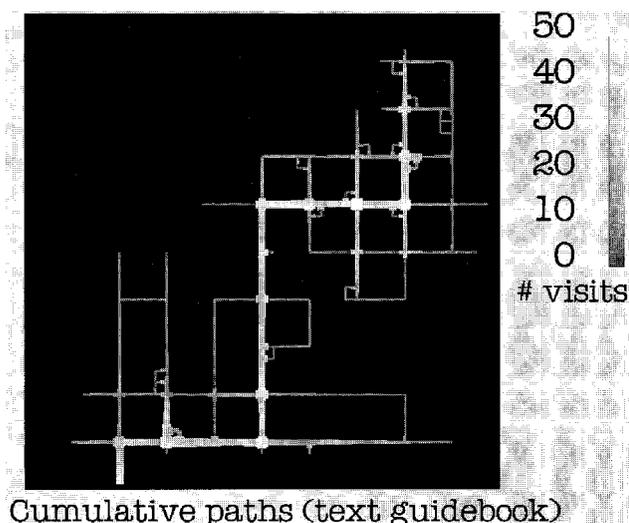
One surprising finding was a high inverse correlation (-0.77 ; $p = 0.01$) between spatial skills and the time spent consulting the text guidebook. Subjects with better spatial ability seemed able to form a mental picture of a landmark based on textual description more quickly than did subjects with poorer

spatial ability. This is not the case for image and worldlet guidebooks, however.

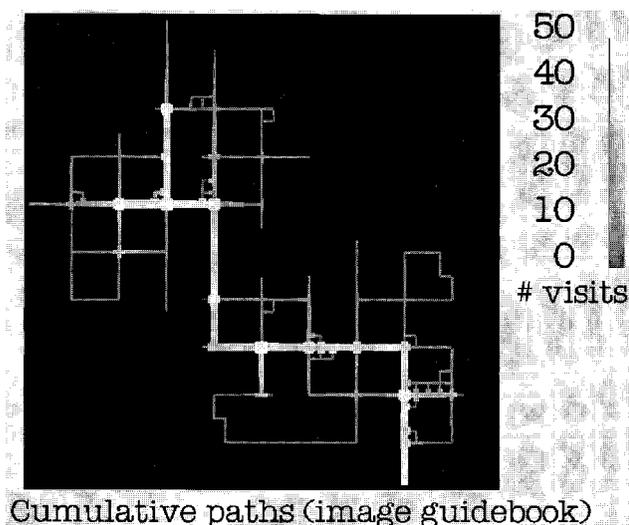
We expected to find a negative correlation between spatial skills and overall travel time in the image and worldlet guidebook cities. The data did not support this, nor did they support another of our hypotheses that there would be a positive correlation between verbal skills and the performance in the text guidebook city. Verbal skills made no difference in overall wayfinding time.

7.5 Observations on Subject Strategies

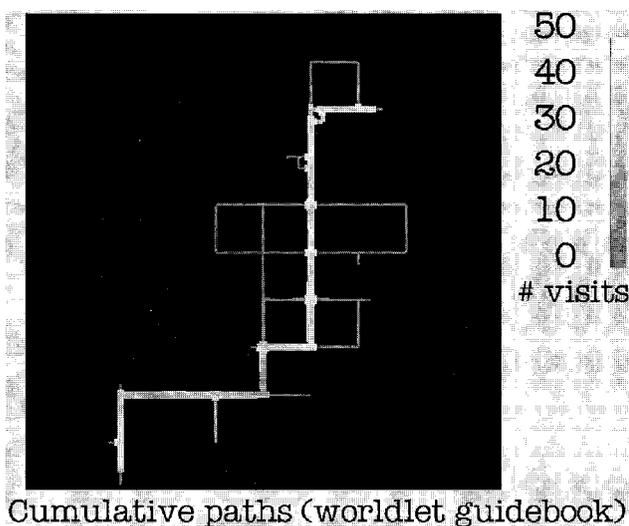
Many subjects said that the tasks were harder than they had imagined, and afterwards several said that they felt tired. No subjects complained of mo-



Cumulative paths (text guidebook)



Cumulative paths (image guidebook)



Cumulative paths (worldlet guidebook)

Table 3. Subjects Were Asked to Rate the Difficulty of the Search Task Using the Three Different Types of Guidebooks

	Text	Image	Worldlet
Very easy	0	2	11
Easy	10	8	9
Doable	10	8	6
Difficult	5	7	0
Very difficult	1	1	0
Median	Doable	Doable	Easy

tion sickness. Most subjects required 70 min. or more to finish all aspects of the computer-based part of the experiment.

We observed subjects to use three general types of strategies while deciding what direction to go from a street intersection:

- Look down each street from many viewpoints within the intersection until a landmark or landmark context is seen.
- Look down each street and, if nothing looks familiar, search exhaustively in each direction until the next landmark is found.
- Look down each street and, if nothing looks familiar, try going one-half block in each direction and look again.

Subjects using the first strategy seemed to find their way more quickly and with less backtracking than did subjects who used the second and third strategies. Another strategy, described by some subjects on their written questionnaire, involved viewing a landmark in the worldlet guidebook from many distant viewpoints in the hopes that one of the viewpoints would be similar to the direction from which they would actually approach the

Figure 13. In these street plots, line thickness and brightness indicates the number of times that a street was traversed. In top and middle, the 26 subjects traversed some streets 47 times indicating backtracking. Top and middle also exhibit circuitous trails. In bottom, very little backtracking occurred, and circuitous trails are minimal.

landmark in the city. We observed that some subjects often forgot what it was they were looking for, and so had to open the guidebook momentarily to remind themselves.

Not surprisingly, landmarks with prominent signage (such as “Bob’s Grocery”) and landmarks painted with more highly saturated colors seemed easier to find than other types of landmarks. Also, after passing all the landmarks, several subjects exclaimed that they had lost all sense of direction within the world.

A final unexpected result were the many similarities between the text and image guidebooks. The text and image guidebooks were only significantly different in the time subjects spent consulting them. Considering that text and images represent landmarks in very different ways, we expected differences in time and distance measures.

7.6 Major Findings

In interpreting these findings, it is important to remember that the subjects performed the wayfinding tasks using only the landmark knowledge gained by studying the guidebooks. No procedural or survey (map) information was provided.

- As we predicted, the total time spent traveling to the goal was significantly less using the worldlet guidebook than when using the text or image guidebooks.
- Subjects spent significantly more time studying the worldlet guidebook than they did the text or image guidebook.
- Increased worldlet study time was more than compensated for by a reduction in the time required to reach the goal. Using the worldlet guidebook, subjects traveled an average of 55% faster than when using the text guidebook and 38% faster than when using the image guidebook.
- Increased worldlet study time also allowed subjects to move more directly to the goal and reduced backtracking time to almost zero. This is shown diagrammatically in figure 13 in which figures 13a and 13b show brighter, wider streets than in figure

13c. Figure 13c also exhibits fewer detours and less wandering than do figures 13a and 13b.

We postulate that, compared to the landmark knowledge gained from text and image guidebooks, the improved landmark knowledge built while consulting the worldlet guidebook enabled subjects to refer less frequently to the guidebook. This postulation is further supported by examining worldlet guidebook results:

- Subjects spent significantly more time studying the worldlet guidebook than they did the other types of guidebooks.
- When using the worldlet guidebook, subjects traveled faster, more directly, and with more skill, to the goal kiosk.
- Subjects’ median rating for the difficulty of the search task was *Easy* when using the worldlet guidebook.

8 Conclusion

The wayfinding literature shows that landmarks are critical in navigating an environment, whether it is real or virtual. While building landmark knowledge, travelers characterize a landmark representation by its 3-D shape, size, texture, context, and so forth. In a wayfinding task, textual or image guidebook landmark descriptions fail to engage the full range of 3-D landmark characteristics recognized and used by travelers to find their way.

We know of no studies that have considered the comparative advantages of different types of guidebooks for wayfinding. Intuitively, one would predict that 3-D thumbnails make better guidebooks for 3-D worlds than do text and 2-D thumbnail images, and, indeed, this is what we found. We wanted, however, to identify the component activities in using guidebooks that most benefitted from 3-D elements. Perhaps of greatest interest was our finding that worldlets enabled subjects to find their way in complex environments with almost no backtracking. By spending more time studying the worldlets in their guidebooks, they spent less time looking around, reorienting themselves, conducting exhaus-

tive searches, and performing similar time-wasting activities. We must be cautious, however, in what we conclude because worldlets contain more information than do text or 2-D thumbnail images. To control for this difference in information content, more experiments are necessary.

9 Future Work

We would like to investigate the comparative advantages of worldlet guidebooks for wayfinding tasks in corresponding virtual and real world environments (both urban and rural). We are also considering an experiment to evaluate tandem use of traditional 2-D maps and landmark guidebooks. Requiring effective map and compass use in a natural environment, the sport of orienteering (Kjellström, 1976) might benefit from worldlets. By encapsulating the region around the flagged control points into a worldlet, orienteers could gain first-person experience with natural landmarks such as foliage, rocks, gullies, shadows, and so on, before beginning the search.

Acknowledgments

Our appreciation goes to Planet 9 Studios (www.planet9.com) for the use of their city building textures. Thanks to John Moreland and Mike Bailey for technical assistance on this project, and thanks to Suzi Feeney for collaborating in the pilot study design. UCSD librarian Debbie Cox and her staff were invaluable in locating references. Thanks also to John Helly and Reagan Moore for their support. San Diego Supercomputer Center is funded by NSF, the University of California, the State of California, and industrial partners.

References

- Appleyard, D. A. (1976). *Planning a Pluralistic City: Conflicting Realities in Ciudad Guayana*, Cambridge, MA: MIT Press.
- Art+Com. (1997). The cybercity Berlin project [on-line]. Available: <http://www.artcom.de/projects>.
- Chen, S. E. (1995). QuickTime VR—An image-based approach to VE navigation. *Proceedings of SIGGRAPH 95*, 29–38.
- Darken, R. P., & Sibert, J. L. (1996). Wayfinding strategies and behaviors in large virtual worlds. *Proceedings of CHI 96*, 142–149.
- Downs, R. J., & Stea, D. (1973). Cognitive maps and spatial behavior. In R. J. Downs & D. Stea (Eds.), *Image and Environment*. Chicago: Aldine Publishing Company, 8–26.
- Elvins, T. T. (1997). Wayfinding in layered information worlds. *Proceedings of the CODATA Euro-American Workshop on Data and Information Visualization*, 6–7.
- Elvins, T. T., Nadeau, D. R., & Kirsh, D. (1997). Worldlets—3D thumbnails for wayfinding in virtual environments. *Proceedings of the ACM User Interface Software and Technology Symposium*, 21–30.
- Elvins, T. T. (1998). *Worldlets—A novel 3D thumbnail affordance for wayfinding in large virtual environments*. Unpublished doctoral dissertation, Department of Electrical and Computer Engineering, University of California—San Diego.
- Elvins, T. T., Nadeau, D. R., Kirsh, D., & Schul, R. (1998). Worldlets—3D thumbnails for 3D browsing. *Proceedings of the ACM CHI Conference*, 21–30.
- Foley, J., van Dam, A., Feiner, S., & Hughes, J. (1990). *Computer Graphics Principles and Practice*. Reading, MA: Addison-Wesley.
- Gillner, S., & Mallot, H. A. (1998). Navigation and acquisition of spatial knowledge in a virtual maze. *Journal of Cognitive Neuroscience*, 10(4), 445–463.
- Goldin, S. E., & Thorndyke, P. W. (1982). Simulating navigation for spatial knowledge acquisition. *Human Factors*, 24(4), 457–471.
- ISO. (1997). *The Virtual Reality Modeling Language*. International Standard ISO/IEC 14772-1:1997 [on-line]. Available: <http://www.vrml.org/Specifications/VRML97/>
- Kjellström, B. (1976). *Be Expert with Map and Compass: The Complete Orienteering Handbook*. New York: Scribner Publishers.
- Liggett, R., Friedman, S., & Jepson, W. (1995). Interactive design/decision making in a virtual urban world: visual simulation and GIS. *Proceedings of the 1995 ESRI User Conference, Environmental Systems Research* [on-line]. Available: <http://www.esri.com/base/common/userconf/proc95>.
- Lynch, K. (1960). *The Image of the City*. Cambridge, MA: MIT Press.

- Passini, R. (1992). *Wayfinding in Architecture* (2nd ed.). New York: Van Nostrand Reinhold.
- Pausch, R., Burnette, T., Brockway, D., & Weiblen, M. E. (1995). Navigation and locomotion in virtual worlds via flight into hand-held miniatures. *Proceedings of SIGGRAPH 95*, 399–400.
- Peponis, J., Zimring, C., & Choi, Y. K. (1990). Finding the building in wayfinding. *Environment and Behavior*, 22(5), 555–590.
- Ruddle, R. A., Payne, S. J., & Jones, D. M. (1998). Navigating large-scale desk-top virtual buildings. *Presence: Teleoperators and Virtual Environments*, 7(2), 179–192.
- Satalich, G. A. (1995). *Navigation and Wayfinding in Virtual Reality: Finding the Proper Tools and Cues to Enhance Navigational Awareness*. Unpublished masters thesis, Department of Computer Science, University of Washington [on-line]. Available: <http://www.hitl.washington.edu/publications/>.
- University College London Center for Advanced Spatial Analysis. Virtual London: A proposal [on-line]. Available: <http://www.casa.ucl.ac.uk/30days/london.htm>.
- Waller, D., Hunt, E., & Knapp, D. (1998). The transfer of spatial knowledge in VE training. *Presence: Teleoperators and Virtual Environments*, 7(2), 129–143.
- Wechsler, D. (1955). *Manual for the Wechsler Adult Intelligence Scale*. New York: Psychological Corporation.
- Wernecke, J. (1994). *The Inventor Mentor*. Reading, MA: Addison Wesley.
- Wickens, C., & Baker, P. (1995). Cognitive issues in virtual reality. In Woodrow Barfield and Thomas Furness III (Eds.), *Virtual Environments and Advanced Interface Design* (pp. 514–541). New York: Oxford University Press.
- Witmer, G. G., Bailey, J. H., Knerr, B. W., & Parsons, K. C. (1996). Virtual spaces and real world places: Transfer of route knowledge. *Int. J. Human-computer Studies*, 45, 413–428.