SPATIAL ORIENTATION, WAYFINDING, AND REPRESENTATION

Rudolph P. Darken and Barry Peterson
Department of Computer Science
Naval Postgraduate School
Monterey, California 93943-5118
831-656-4072
831-656-4083 (fax)
darken@acm.org
peterson@cs.nps.navy.mil

Darken, R.P., & Peterson, B. (2001). Spatial Orientation, Wayfinding, and Representation. Handbook of Virtual Environment Technology. Stanney, K. Ed.

1. INTRODUCTION

Everyone has been disoriented at one time or another. It is an uncomfortable, unsettling feeling to be unfamiliar with your immediate surroundings and unable to determine how to correct the situation. Accordingly, we might think that the goal of navigation research in virtual environments is to create a situation where everyone is oriented properly all the time and knows exactly where everything is and how to get there. This, however may not be absolutely correct. Much is gained from the navigation process beyond just spatial knowledge. The path of discovery rarely lies on a known road. The experience of serendipitous discovery is an important part of human navigation and should be preserved. But how do we resolve the conflicts between this and the not-so-pleasant experience of lostness?

Navigation tasks are essential to any environment that demands movement over large spaces. However, navigation is rarely if ever the primary task. It just tends to get in the way of what you really want to do. Our goal is to make the execution of navigation tasks as transparent and trivial as possible, but not to preclude the elements of exploration and discovery. Disoriented people are anxious, uncomfortable, and generally unhappy. If these conditions can be avoided, exploration and discovery can take place.

This chapter is about navigation in virtual environments -- understanding how people navigate and how this affects the design of virtual environment applications. We begin with a clarification of terms and some theoretical background on navigation, primarily in the real world. A discussion of methods for navigation performance enhancement will follow. This is about how to improve performance in a virtual environment. This is different from the next topic concerning the use of virtual environments as training tools where we are interested in performance in the real world. Lastly, the chapter concludes with a summary of principles for the design of navigable virtual environments.

1.1 Definition of Terms

One of the problems we find in the literature is confusion over terms involving navigation. It is difficult to compare two studies that use different terms in different ways. We may be unknowingly comparing apples to oranges. We use specific terms with specific definitions and encourage the research community to adopt these.

Wayfinding is the cognitive element of navigation. It does not involve movement of any kind but only the tactical and strategic parts that guide movement. As we will see later in the chapter, wayfinding is not merely a planning stage that precedes motion. Wayfinding and motion are intimately tied together in a complex negotiation that is navigation. An essential part of wayfinding is the development and use of a *cognitive map*, also referred to as a *mental map*. Still poorly understood, a cognitive map is a mental representation of an environment. It has been called a "picture in the head" although there is significant evidence that it is not purely based on imagery but rather has a symbolic quality. The representation of spatial knowledge in human memory that constitutes a cognitive map will be an important part of this chapter.

Motion is the motoric element of navigation. A reasonable synonym for motion is *travel* as used by Bowman, et al. (1997). Durlach and Mavor (1995) subdivide this further into *passive transport* such as a "point-and-fly" interface or other abstraction, and *active transport* such as the Omni-Directional Treadmill or other literal motion interface (Darken, Cockayne & Carmein, 1997). Active transport interfaces are often referred to as *locomotion* interfaces. *Maneuvering* is a subset of motion involving smaller movements that may not necessarily be a part of getting from "here" to "there" but rather adjusting the orientation of perspective, as in rotating the body, or sidestepping. This is an important distinction to make for the development of active transport interfaces for locomotion such as Gaiter (Templeman, Denbrook, & Sibert, 2000). (See also Chapter 10).

Navigation is the aggregate task of wayfinding and motion. It inherently must have both the cognitive element (wayfinding), and the motoric element (motion). Consequently, we use this term only when we mean to imply the aggregate task and not merely a part. The literature is replete with references to "navigation" that are only interested in novel motion techniques. We find this to be confusing and counterproductive to the discussion.

We should also mention what is implied by *navigation performance*, as this is the metric we need to use to determine the relative effectiveness of specific navigation tools and techniques. This is entirely dependent on the navigation task in question. Are we studying the ability of a person to find an unknown object in a complex space? – then search time might be an appropriate measure. Are we interested in the ability to find a known location in a complex space? – then route following might be appropriate. Are we interested in a person's overall knowledge of the configuration of a space? – then a map drawing exercise might be appropriate. These issues will be discussed in more detail in section 2.

1.2 Training Transfer or Performance Enhancement?

There are two primary classes of applications having to do with navigation in virtual environments. All virtual environments that simulate a large volume of space will have navigation problems of one sort or another. Typically, any space that cannot be viewed from a single vantage point will exhibit these problems as users move from one location to another. The need to maintain a concept of the space and the relative locations between objects and places is essential to navigation. This is called *spatial comprehension*, and like verbal comprehension, involves the ability to perceive, understand, remember, and recall for future use.

In applications where we find that users tend to become disoriented or are unable to relocate previously visited points of interest, it is desirable to either redesign these applications so these problems do not appear or provide tools or mediators to help alleviate these problems. This class of application involves a need to enhance performance within the virtual environment. This distinction differentiates these applications from those where improved performance is required outside of the virtual environment, in the real world. These are training transfer applications which we will discuss next.

The second class of applications that involves navigation has to do with the use of virtual environments as training aids for real world navigation tasks. The fact that we can construct virtual representations of real environments has lead many to consider the use of virtual environments for much the same purposes that a conventional paper map might be used. While there are certainly similarities between these two classes of applications, the validation process is entirely different. If we want to show that a visualization scheme in a virtual building walkthrough, for example, can be used to lessen cases of disorientation, we need only show that users of the virtual environment with this visualization scheme perform better on navigation tasks than users of an identical virtual environment without the visualization scheme. However, if we want to show that this same virtual environment can be used as a training aid for navigation tasks in the actual building, navigation performance within the virtual environment, while interesting, does not prove our point. A training transfer study must be completed to show that users who trained on the virtual environment navigate the building better than users who received some other form of training or possibly no training at all. This is an example of the use of a virtual environment as a training aid for specific environments. Navigation performance is expected to improve in one specific real environment, and nowhere else.

Another interesting issue is the use of a virtual environment as a training aid for general navigation tasks. If it could be shown that a virtual environment could help people use paper maps more efficiently, or to select landmarks in an environment effectively, that performance increase would be expected to exist across physical spaces, possibly assuming some spatial similarities[*] with that of the training environments. This is beyond the scope of our discussion here but is an important use of virtual environments that has not been explored as yet.

2. BACKGROUND

Whether or not a virtual environment attempts to simulate the real world, we have to consider the fact that people are accustomed to navigation in physical spaces. Certainly, there are differences between regions of physical space that alter how navigation works, such as navigating in a forest versus navigating in a city, but there are assumptions that can be made based on past experience in real environments that are useful in any real space. While

[*] For example, a virtual environment that could be used to train a person to effectively navigate in a generic city would probably have little impact on that person's ability to navigate in a generic forest.

this is not a satisfactory reason to blindly copy the real world in every way we can, we certainly have to learn everything possible about how people relate to the physical world so we can understand how to build better virtual environments.

2.1 Spatial Knowledge Acquisition

There are many ways to acquire spatial knowledge of any environment. The fundamental distinction between sources of spatial knowledge is whether the information comes directly from the environment (primary) or from some other source (secondary) such as a map. An issue specific to secondary sources has to do with whether or not the source is used inside or outside of the actual environment.

2.1.1 Direct Environmental Exposure

When you navigate in an environment, you extract information both for use on whatever navigation task currently is being executed and for any subsequent navigation task. Exactly what information is useful for navigation? We can't possibly attend to every stimulus and make use of it. Much of it is irrelevant or at least of lesser importance. If we knew what information was the most important, this could be useful in designing virtual environments. Designers would know what to put where to help people find their way around. While there is no clear answer to this question, Kevin Lynch presents the most compelling, environment independent answer to date (Lynch, 1960).

In studying urban environments (Lynch was an urban planner), he found that there are certain similarities that cross cities. There are in fact "elements" of urban environments; building blocks that can be used to construct or decompose any city. He starts with *landmarks* which are the most salient cues in any environment. They are also directional meaning that a particular building might be a landmark from one side but not another. Then there are *routes* (or *paths* if you prefer) that connect the landmarks. They don't necessarily connect them directly but they move you through the city such that the spatial relationships between landmarks become known. A route isn't necessarily a road. It could also be a bike path or railroad. Cities tend to have complex road and rail structures. Interchanges or junctions between routes are called *nodes*. These are important because they are fundamental to the structure of the routes. This structure must be understood before proficient navigation can take place. Most cities have specific regions that are explicitly or implicitly separated from the rest of the city. These are *districts*. Landmarks and nodes typically live in districts. Routes pass through districts and connect them. Finally, regions of the city, and in fact the city itself, are bounded by *edges*. Edges prevent or deter travel. A typical edge is a river or lake.

An interesting fact is that classifying some city object as one element or another does not preclude it being classified as something else in another context. To a pedestrian, a walking path is a route while highways and railroads are edges. To a driver, the roads are routes and everything else is an edge. It depends on the mode of travel. Furthermore, the mode of travel also effects *what* gets encoded, not just *how* it gets encoded (Goldin & Thorndyke, 1982).

2.1.2 Map Usage

There are a number of secondary sources that have been used for spatial knowledge acquisition. These include maps, photographs, videotape, verbal directions, and recently virtual environments. The most common of these is the map. For any map used in any environment, virtual or real, we need to know when the map is used, or more appropriately, what tasks the map is to be used for. The critical issue is in whether or not the map is to be used preceding navigation or concurrent with navigation. This is important because maps that are used concurrently with navigation involve the placement of oneself on the map. The first part of any task of this nature is, "Where am I? What direction am I facing?". A transformation is required from the egocentric perspective to the geocentric perspective. If the map is used as a precursor to navigation, it is used only for planning and familiarization. No perspective transformation is required. Planning a trip is one example of such a geocentric navigation task. The planning is done outside of the environment so there is no perspective transformation needed. However, when such a transformation is required, as in using a map during navigation, the problem is more complicated.

When a perspective transformation must be performed, the rotation of the map can have a great effect on performance. Aretz and Wickens (1992) showed that maps used during navigation tasks (egocentric tasks) should be oriented "forward-up" (the top of the map shows the environment in front of the viewer) while maps used for planning or other geocentric tasks should be "north-up". This same concept is reinforced by Rossano and Warren (1989) who showed that judgements of direction are adversely affected by misaligned maps. Levine (1984) goes on to provide basic principles of map presentation including aligning the map with the environment, always showing two concurrent points on the map so the viewer can triangulate position, and avoiding symmetry. Péruch, Pailhous, and Deutsch (1986) showed that redundancy of information is crucial to resolve conflicts between different frames of reference.

The key to map use for navigation is resolving the egocentric to geocentric perspective transformation. This certainly isn't the entire problem, but it is the biggest part of it. This involves the ability to perform a mental rotation. The easier this rotation is, the easier the task is. Unfortunately, mental rotation is not a level playing field. Some of us are better at it than others, so much so that it affects the way individuals perform navigation tasks (McGee, 1979; Thorndyke & Stasz, 1980). We all know individuals who maintain a high level of proficiency in navigating environments, including environments they have not ever been in before. Alternatively, there are individuals who will get lost on the way home from work if they can't follow their usual route for some reason.

One thing we might want to do to simplify the transformation is to show the viewpoint position and orientation directly on the map. Until recent improvements in geographic positioning system (GPS) technologies, this was not possible to do in the real world. As of May 1, 2000, the United States no longer jams GPS signals for private use so GPS is now 10 times more accurate than is was previously. While "dithered" GPS could only place a receiver in an area the size of a football field, now it is accurate to within the area of a tennis court. However, there are no commercially available databases available that make use of spatial data at this resolution, and even if there were, we still do not know what the best uses are for this data in the real world, so it will be some time before we realize a payoff from the declassification of GPS. Automobile navigation systems with moving maps are already commonplace. However, the interfaces to these systems are generally poor. One of the issues we will look at later is the use of virtual maps and how they should be presented.

2.2 Representations of Spatial Knowledge

Arguably the most important part of this puzzle is the part we understand the least. Once spatial knowledge is acquired, how is it organized in the mind for future use? Spatial knowledge must be organized in some way such that it can be used during navigation tasks. The term "cognitive map" was first used by Tolman to describe a mental representation of spatial information used for navigation (Tolman, 1948). However, fifty years later, we still don't have any hard answers about the structure of spatial knowledge.

The representation of spatial knowledge is affected by the method used to acquire it. Knowledge acquired from direct navigation is different from knowledge acquired from maps. After studying a map of San Francisco before an initial visit there, information gleaned from the map is in a north-up orientation, just like the map. Consequently, when entering the city from the north, the mental representation of the city that has been acquired is upside down from what is being seen. This requires a 180° mental rotation before the cognitive map can be in line with the real environment. Spatial knowledge acquired from maps tends to be orientation specific (McNamara, Ratcliff & McKoon, 1984; Presson, DeLange & Hazelrigg, 1989; Presson & Hazelrigg, 1984). This implies that accessing information from a misaligned cognitive map is more difficult and error prone than if it were aligned (Boer, 1991; Rieser, 1989).

The longest standing model of spatial knowledge representation is the Landmark, Route, Survey (or LRS) model described by Seigel and White (1975) and Thorndyke and Goldin (1983). This model not only addresses spatial knowledge but also the development process. The theory states that we first extract landmarks from an environment. These are salient cues but are static, orientation dependent, and disconnected from one another. Landmark knowledge is like viewing a series of photographs. Later, route knowledge develops as landmarks are connected by paths. These need not be optimal paths. At this point, I may know how to get from A to B and from B to C but I probably don't know a direct route from A to C. Route knowledge can be thought of as a graph of nodes and edges that is constantly growing as more nodes and edges are added. Finally, survey knowledge (or configurational knowledge) develops as the graph becomes complete. At this point, even if I have not traversed every path through my environment, I can generate a path on-the-fly since I have the ability to estimate relative distances and directions between any two points. This model directly fits the elements of urban environments described by Lynch (1960).

The most important caveat to this development process has to do with the use of maps. Maps allow us to jump over the route knowledge level and proceed directly to survey knowledge since they afford a picture of the completed graph all at one time. However, there is no free lunch. Survey knowledge attained from maps is inferior to survey knowledge developed from route knowledge and direct navigation because of the orientation specificity issue described earlier.

A second model of spatial knowledge is similar to the LRS model but it is hierarchical (Stevens & Coupe, 1978; Colle & Reid, 1998). In some cases, direct exposure to an environment for extremely long durations never results in survey knowledge (Chase, 1983). In other cases, survey knowledge develops almost immediately. The model proposed by Colle and Reid suggests a dual-mode whereby survey knowledge can be acquired quickly for local regions and slowly for remote regions. The "room effect" comes from the ability to develop survey knowledge of a room rather quickly but survey knowledge of a series of rooms develops relatively slowly and with more errors.

2.3 Models of Navigation

Understanding how navigation tasks are constructed is useful in determining how best to improve performance. If it were possible to decompose navigation tasks in a general way, we might be able to determine where assistance is needed, or where training can occur. Several attempts have been made at such a model (Chen & Stanney, 2000; Darken, 1996; Downs & Stea, 1977; Neisser, 1976; Passini, 1984; Spence, 1998) but most are either too specific to one type of environment or another or they do not capture the intricacies of the entire task.

The model proposed by Jul and Furnas (1997) is relatively complete in that it incorporates the motion component into the process in a way not attempted before (See Figure 1). The model works like this. I'm at the shopping mall and decide I need a pair of shoes. I have just formulated a goal. Now, how should I go about finding shoes? I decide to try the department stores. Department stores are typically on the far points of the mall. I have just formulated a strategy. The next step is to gather information so I don't walk off in a random direction. I decide to seek out a map of the mall. I am acquiring information and scanning (perceiving) my environment. This is the tight wayfinding/motion loop referred to earlier. I view my surroundings, assess my progress towards my goal, and make judgements as to how to guide my movement. At any time in this loop, I may decide to stop looking for shoes and look for books instead. This is a change in goal. I could also decide to look for a small shoe store instead of a department store. This is a change of strategy. In any case, the task continues, shifting focus and process as necessary.

An important point is that navigation is a situated action (Suchman, 1987). Planning and task execution are not serial events but rather are intertwined in the context of the situation. It is not possible nor practical to consider the task, the environment, and the navigator as separate from each other. Observable and measurable behavior is a product of these factors, yet the relationships between them are, as yet, poorly understood.

In the real world, this process is performed so often that it is typically automatic. When we know where we want to go, we go there. When we don't know where to go, we ask someone or look for some other source of assistance. Virtual environments are an entirely different animal. Movement isn't typically so easy. We have to think about it. Knowing where to go is problematic. Who are you going to ask for help? How do you get un-lost? These are the issues we will discuss next.

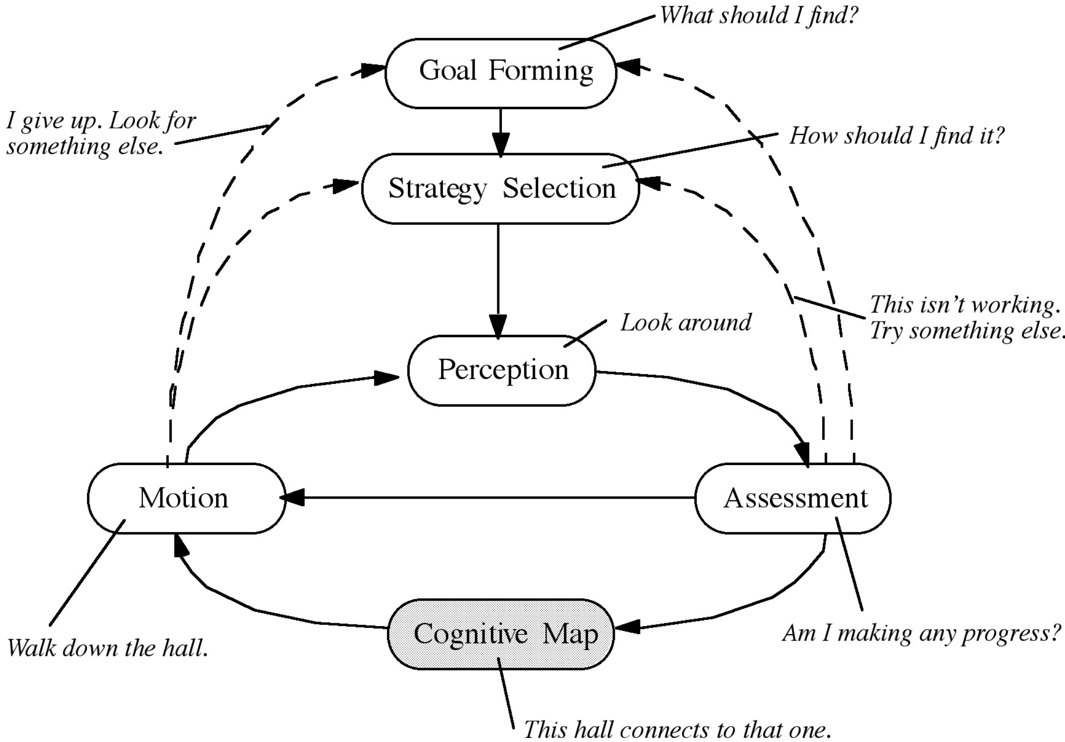

FIGURE 1. A model of navigation adapted from Jul and Furnas (1997).

3. NAVIGATION PERFORMANCE ENHANCEMENT

Many of the physical environments that the Navy must operate in are inherently sparse and consequently difficult to navigate. Literal virtual representations of these environments are even harder to navigate considering that the number of cues available is diminished even further. Early prototypes taught us valuable lessons in how difficult navigation can be. While our open ocean Naval applications might be on the more difficult end of the complexity scale, these same problems can arise in any application. This section will look at ways to help people navigate a virtual space, without regard to whether or not their improved performance transfers outside of the virtual environment.

The most obvious way to address navigation problems in complex virtual environments is to provide some sort of tool or mediator that can be used directly on the task at hand. This, of course, is the history of the map, the compass, the sextant, and the chronometer, to name a few real world tools. Used together, these tools are able to help their users determine their position in the environment, their direction of travel, and the relative position of other objects or places in the environment.

An alternative to mediators is the actual organization of the space itself. There are a few real world vocations related to this including architectural design and urban planning, each of which has extensively studied the relationship between people and their environment. They are interested in much more than just navigation, of course, but there is much to learn from these disciplines about how to construct space in a meaningful way in which people can comprehend and operate effectively. In virtual environments where the contents of the environment cannot be changed, such as in a flight simulator, there isn't much that can be done in the way of spatial organization. However, in other applications, such as many scientific visualizations, it is possible to organize the data expressly for the purpose of navigability.

3.1 Navigation Tools and Mediators

This section discusses a number of navigation tools and mediators that we have investigated in our laboratory. While they are similar to others in the literature, it should be noted that our intent is to investigate principles for the design of navigable virtual environments, not merely new techniques. It is unlikely that one of these tools as described here will perfectly match the needs of a real application. But by explaining how certain types of information affect behavior and its related performance, designers can mix and match the techniques described here and elsewhere to construct custom built tools specifically for the demands of their application.

3.1.1 Maps

The navigation mediator people are most familiar with is the map. Maps are extremely powerful tools for navigation because of the wealth of information they can provide and the rate at which people can digest this information. However, their use in any virtual environment application is not to be taken lightly. There are right ways and wrong ways to use maps. Maps come in a variety of forms, the differences usually being in terms of symbology or projection. However, virtual environments have certain qualities the real world doesn't have that make the use of maps in virtual environments different.

It is possible to navigate directly on a map in a virtual environment. Rather than use the map to determine where to go in the virtual environment, why not just point on the map to where you want to go? This has been attempted several times. The Worlds-In-Miniature (Stoakley, Conway & Pausch, 1995) metaphor was one such implementation. In this case, a scaled-down version of the world, a virtual map, was held in the hand. Movement could be specified directly on the virtual map. This same behavior is possible in some video games. DOOM™[†] allows movement to take place while viewing the map. The environment cannot be seen during this interaction. Similarly, the use of maps can be a moded or unmoded task in a virtual environment, meaning that its use can be in lieu of, or concurrent with motion. In the real world, it is typically unmoded. Under certain conditions, you might want to stop moving to read a map but it isn't required. Some games and virtual environments mode the map so that map use precludes motion.

For maps of very large virtual environments, there is a scaling problem. How do I view the map such that I can see the detail I need to navigate but still maintain a sense of the overall space. This is a classic problem of navigation in any problem domain (e.g. Donelson, 1978; Furnas, 1986). There are ways to zoom into a map or otherwise scale it to a usable level. One advantage of virtual environments is that often, the problem only concerns maintaining one perspective -- yours. In applications where I might need to not only know where I am but other people as well, the problem gets considerably harder. Street Atlas USA™[‡] is an excellent example of an application addressing this issue. As you zoom into the atlas to the street level, context is lost. Consequently, there is always an overview window available.

The last issue we will discuss about maps, and the primary focus of this section has to do with the orientation of the map. As we discussed earlier, the orientation of a map with respect to the viewer has a strong effect on the viewer's ability to perform the mental rotation required to use a map during navigation. However, if we already know that forward-up maps are best for egocentric tasks and north-up maps are best for geocentric tasks, isn't this a moot point? The fact that most video games continue to use north-up maps is enough to warrant our investigation. Is it possible that we can put enough redundant information on a north-up map to make it equivalent to a forward-up map?

To start out, we created two very large virtual environments, one was a sparse open ocean environment and the other was an urban environment (Darken & Cevik, 1999). We used both a head-mounted display and a Fakespace PUSH™ display with similar results. The results reported here are using the PUSH™ display. Participants were asked to locate objects in the space. Sometimes they were shown the targets on the map (a targeted search), sometimes they had to return to known targets that were not shown on the map (a primed search), and other

---

[†] Id Software, Inc.
[‡] DeLorme, Inc.

times they had to locate a target not shown on the map and not seen before (a naïve search). We timed their performance and marked wrong turns where participants clearly moved away from a target rather than towards it. At the conclusion of each trial, participants were asked to sketch a map of the environment from memory.

We created two virtual maps; one was in a north-up configuration (See Figure 2), the other was in a forward-up configuration (See Figure 3). Both maps had a "you-are-here" marker that dynamically moved across the map as the user moved through the environment. The only difference on the maps, besides their orientation, was that the you-are-here indicator on the north-up map was a cone while the you-are-here indicator on the forward-up map was a small sphere. This was necessary because a north-up map does not indicate direction implicitly while the forward-up map does.

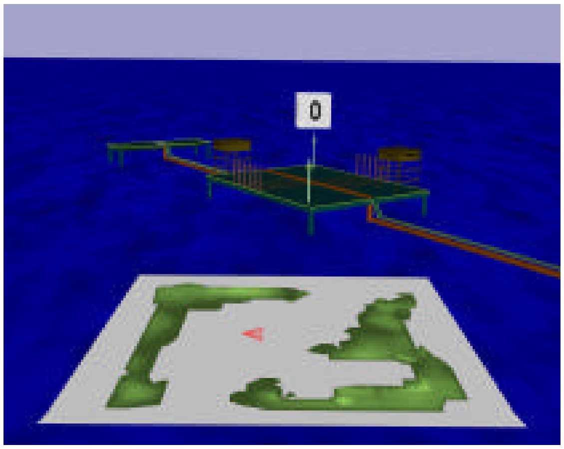

FIGURE 2. The north-up map configuration has a directional cone for the you-are-here indicator. The top of the map always stays at the top. It is aligned with the viewer, not with the environment.

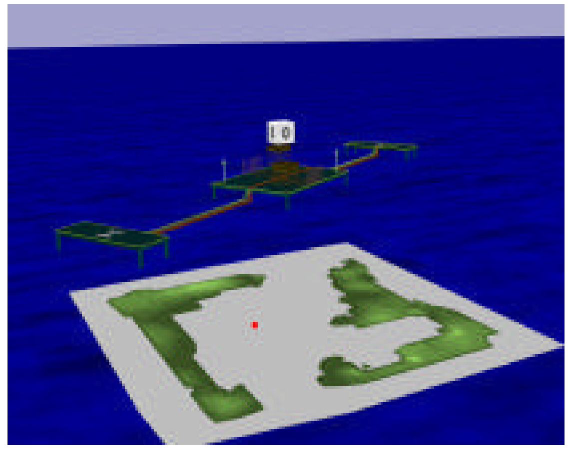

FIGURE 3. The forward-up map configuration has a small sphere in place of the cone for a you-are-here indicator. This map rotates with respect to the viewer as the viewpoint is rotated in the environment. It is aligned to the environment, not the viewer.

What we found was comparable to the results of Aretz and Wickens (1992) but to a lesser degree. The forward-up configuration seems to indeed be best for egocentric navigation tasks while the north-up map is best for geocentric tasks. We also found that individuals with high spatial abilities (as measured with the Guilford-Zimmerman standardized tests) are able to use either type of map better than participants with low spatial abilities on similar tasks. Furthermore, we found that these principles apply across different types of environments with vastly different spatial characteristics, but sparse environments seem to exhibit less of a performance difference than dense environments. Virtual environment designers should make virtual map decisions by carefully weighing the priorities of navigation tasks versus the spatial ability of their users.

Interestingly, we found that in asking participant's subjective assessment, most preferred the north-up configuration even if their performance did not correlate. In post-evaluation debriefing, we learned that most of these people maintain a strong bias based on video game play. Irregardless of how they perform, they like what they're comfortable with. This complicates the problem. Our study seemed to indicate that while the forward-up map was better for egocentric tasks, it was not so much better that we would advise ignoring the preferences of users. It will be imperative in the future to examine the effects of exposure to video games on performance of navigation tasks in virtual environments.

3.1.2 Landmarks

We learned earlier that landmarks are extremely important to spatial knowledge acquisition and representation. As such, it might be useful to allow users to affect the placement of landmarks themselves. If landmarks are useful as "anchors" on which to relatively place other objects, what would happen if the locations of the objects remained constant but we allowed users to insert highly salient landmarks on which to anchor their locations?

Using the open ocean environment as before, participants were given a set of ten different colored markers to place on the surface of the environment at any position. These markers were visible from a far greater distance than the objects themselves. We wondered where they might be placed and how this might affect performance.

Markers tended to be placed in-between objects such that as one marker would disappear from view, another would appear on the horizon. Moving between markers and objects, the participant could always have something to "hold on to" much like a "handrail"[§] is used in sport orienteering. We learned that for most individuals, being in a void, even a partial void like this environment, is very uncomfortable. Most people need regular reassurance that they are not lost. Only the most advanced navigators do not seem to need this kind of assurance.

We also tried this same condition but combined it with a forward-up map (See Figure 4). As participants would place a marker, it would appear both in the environment and on the map. Would the strategy remain the same or would the map override the utility of the markers?

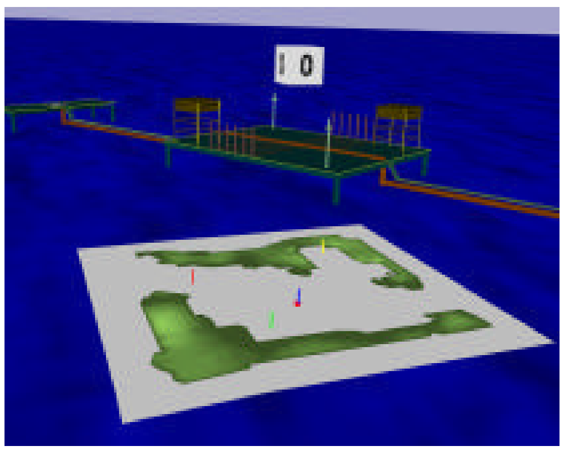

FIGURE 4. This map shows a user returning from finding a number of objects. The markers are used like push-pins to show where targets were located. Before leaving the virtual world and beginning the map drawing task, the participant would simply remember the configuration of markers on the map.

This turned out to be unquestioningly the favorite condition of everything we have tried thus far. The reason is simple. We still didn't give them any help in locating the target objects. But as they found them, they no longer had to remember anything about them. They simply placed a marker at each object. The markers in the environment were not used at all. But the markers on the map were like colored "push-pins". All they had to do was remember what color coincided with what target and the post-trial map drawing exercise was trivial.

Certainly we could devise other techniques that may fare even better than this condition, but it is our intent to study methods that can generalize to many navigation tasks in any type of environment rather than only to the specific tasks and environments used in our studies. The objective is to determine design principles so that application designers have a starting point from which to develop specific mediators for their specific tasks in their specific environments.

[§] A handrail is a linear feature, such as a stream or road, that land navigators use to guide movement. They typically use it to constrain movement, keeping it to one side and traveling along it for some distance.

3.1.3 Trails

Since participants used the markers to create a sort of “trail” connecting objects with markers and other objects, we decided to revisit an old idea. If a trail was left behind such that it specified not only where the participant had been but what direction they were going at that time, this might be an even better tool than the markers (See Figure 5). We called this the “Hansel and Gretel” technique (Darken & Sibert, 1993). A better analogy is that of footprints since footprints are directional and breadcrumbs are not.

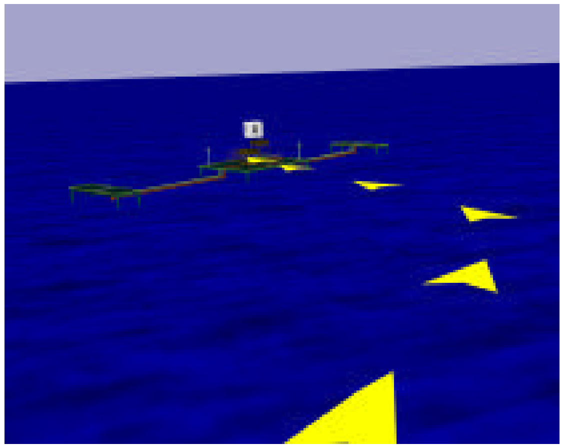

FIGURE 5. In this example, the participant has moved backwards and is looking towards the start object location. The footprints show the direction the user was looking at that point in time.

This technique is useful for scanning space in an exhaustive search (e.g. a naïve search). One of the problems in an exhaustive search is knowing if you’ve been in some place before. An optimal exhaustive search never revisits the same place twice, assuming the targets aren’t moving (ours weren’t). However, we thought it might be useful to retrace steps if disorientation occurred. This was not the case. As navigation proceeds, the environment becomes cluttered with footprints, so much so that the directional component is ignored. Also, when the user crosses paths, how do you disambiguate which ingress route goes with which egress route?

As with the markers, we tried this with a forward-up map (See Figure 6). Similarly, we found that participants ignored the footprints in the environment only using the trail left on the map to help direct the search.

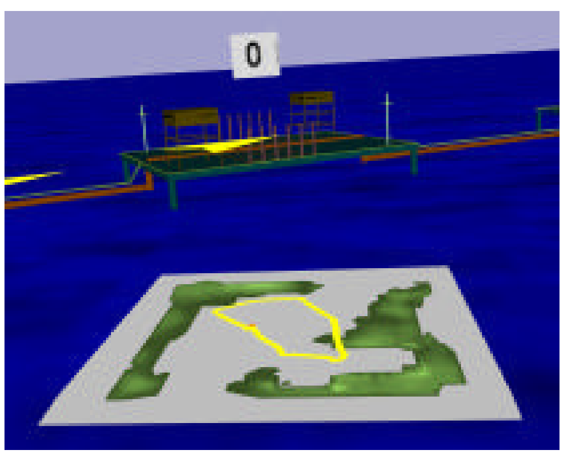

FIGURE 6. This participant has made a cycle through the environment. The actual positions of the target objects are not shown on the map. The participant must remember where along the path they were located.

3.1.4 Direction Finding

So far, we have discussed a number of tools and mediators that deal with absolute position information. The maps showed exactly where the user was at all times. The markers specify an exact location in the environment. Each footprint designates an exact location as well. What about orientation? How important is it to know absolute direction versus absolute position?

To study this, we came up with two simple tools, both of which show nothing more than direction. Even though our environments didn't have a "north" we decided to make the initial view direction be north. It didn't matter what we said north was as long as we were consistent. The compass in Figure 7 always points to virtual north. It floats out in front of the viewpoint similarly to the maps we discussed earlier.

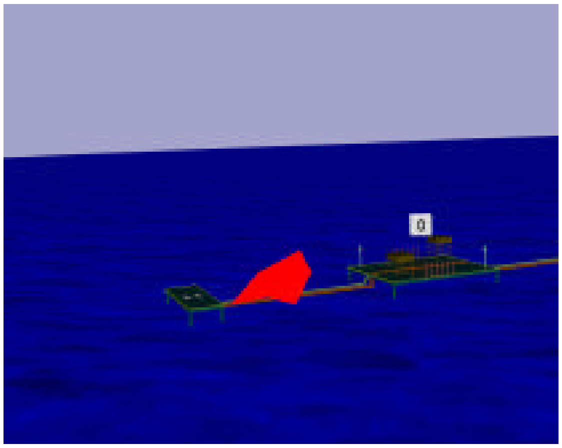

FIGURE 7. The compass simply shows the direction that the participant was facing at the onset of the trial.

The second tool we implemented was a virtual sun. This was placed on the horizon such that it would identify one direction only (See Figure 8). It was up to the participant to call it east or west. This never seemed to be a problem.

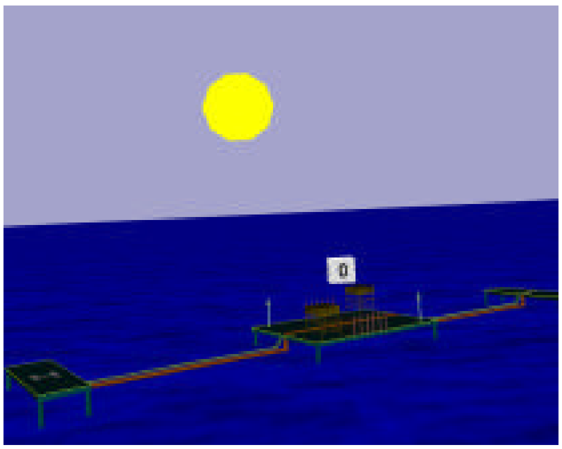

FIGURE 8. The sun is positioned low on the horizon so directional information will be clear and unambiguous.

The task was the same as it was for every other condition. We found that the compass was preferable to the sun because it was in view all the time versus the sun that required the participant to turn towards it. However, performance using both of these tools was very low as compared to others. In fact, performance on both of these conditions was not significantly better than the control condition where no assistance was given whatsoever. The reason for this is that while maintaining spatial orientation is essential to all navigation tasks, orientation without position information is not useful. Similarly, position information without orientation information is equally limited in utility as shown by the "coordinate feedback" tool described by Darken and Sibert (1993).

3.2 Organizational Remedies

It could be said that the tools and mediators described in the previous section should be the last recourse of the designer if all else fails. That is, if the constraints of your application are such that you simply can't develop a navigable virtual environment that doesn't need tools or mediators for usability, then you need to determine what tools are most appropriate to allow any user to navigate effectively. However, there are many applications where it is not appropriate or even possible to organize a space because the positions of objects in the space are constrained in some way. For example, if we are constructing a virtual environment driving simulation of metropolitan Los Angeles, we would have no control over the positions and orientations of buildings and streets. If, alternatively, we were constructing a virtual environment visualization of the stock market, we would have control over how objects looked, where they were located, and how they were organized.

This section will discuss two primary methods for organizing spaces for navigability. The first relies heavily on Lynch's "elements" of the urban landscape where the environment is implicitly organized (Lynch, 1960). The other involves a more explicit organizational technique where a specific pattern is forced on the space itself.

3.2.1 Implicit Sectioning

Passini (1984) talks about the use of an "organizational principle" in architectural design. If a space has an understandable structure and that structure can be made known to the navigator, it will have a great influence on the strategies employed and resulting performance on navigation tasks. For example, knowing that Manhattan is generally a rectangular grid is of great benefit to any navigator. Given the knowledge of the grid and its orientation, words like "uptown" and "downtown" instantly have meaning. But while there is great power in using such an organizational principle, there is a danger that goes along with it. Violations to that principle will have a much greater negative effect than they might otherwise have. For example, we know that Manhattan is generally a grid, but Broadway cuts through on an angle, violating the grid principle. A naïve tourist thinking that the grid principle held throughout, will probably be mislead at some point. It is important to develop a clear organizational principle and stick to it throughout the environment. If it must be violated, it should be made clear where the violation occurs and that it is a violation, otherwise the navigator may attempt to erroneously fit it into a cognitive representation.

Organizational principles can also add meaning to cues that might be seen during navigation. Again using our Manhattan example, if we are looking for 57th St., street signs for 44th St., 45th St., etc. tell us that we are not only going in the right direction but they also give us a rough estimate of distance as well.

There have been attempts to apply Lynch's elements of the urban landscape to virtual environments (Darken, 1996; Ingram & Benford, 1995). In principle, this will work, but implementations can seem contrived as if we are forcing an inappropriate structure on abstract data. It may be that the very generic nature of virtual environments might demand a different set of elements similar to Lynch's but not necessarily identical.

Darken and Sibert (1996) commented that people inherently dislike a lack of structure. A person who is in an environment that is nearly void of useful cues and that does not suggest ways to move through it is generally uncomfortable. Users will grasp at anything they can view as structure whether or not the designer intended it to be that way. In earlier experiments, we observed participants using coastlines and the edges of the world as paths even though this was not a particularly effective strategy towards completion of the task. People do this because it's all they have to work with. If even a simple bit of structure is added such as a rectangular or radial grid (described in the next section) performance immediately improves. A path should suggest to the navigator that it leads to somewhere interesting and useful.

It is common to see urban metaphors applied to unstructured environments such as web sites because the structure of the city can add meaning to information that may be viewed as otherwise amorphous. This works in many cases, even when there is no obvious semantic connection between the information presented and the city metaphor. In other cases, this fails because the constraints implied by an urban landscape do not coincide with the constraints of the information presented. This should be used judiciously with an understanding that again, there is no free lunch. Using a metaphor to simplify navigation may unknowingly inhibit some other piece of functionality. Therefore, it is important when studying user performance in these types of applications to study the whole task, not just the navigation component. Again, navigation is not an isolated action but is situated in some other higher order task.

3.2.2 Explicit Sectioning

In some cases, like our open ocean environment, there are simply too few objects in the environment to organize into a navigable environment. In these cases, the space can be explicitly organized using explicit cues. We experimented with two sectioning schemes based on Lynch's ideas of "districts" and "edges" and also using a simple organizational principle as described by Passini.

The first scheme we tried was a radial grid. We placed a radial grid over our open ocean environment. There were highly salient landmarks in each cardinal direction and a single landmark in the center. The second scheme we tried was basically the same idea but was a rectangular grid rather than radial. Again, highly salient landmarks were placed in the cardinal directions and the center.

We found that search strategies for these two schemes were quite different but that performance between the two was not significantly different. The radial grid tended to elicit strategies having to do with "pie slices" or similar ideas[**], while the rectangular grid produced more "back and forth" motions within bounded regions. The nature of the task or environment would probably dictate which was better in a particular situation. In our case, the difficulty of the task without any aids caused performance to increase for both grids. But observed strategies were clearly related to the type of grid we presented. In a given situation, if a particular strategy is clearly preferable over alternatives, structure such as this can be used to guide the user towards the preferred method. As we observed in our experiments, there will almost always be room for unexpected behavior, but it is possible, given a well defined task, to design-in a desired tendency.

We then experimented with combining the grid overlays with the map tools previously described (See Figure 9). This was found to be highly effective, particularly with the forward-up map because the map provided directional information and survey knowledge while the grid provided directional cues and landmark knowledge. Some participants tended to focus more on the map than the grid or vice versa but in general, the two were used in combination as intended.

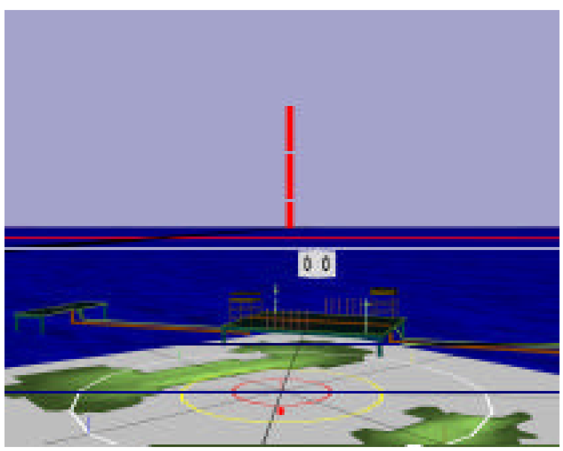

FIGURE 9. The radial grid superimposed on a forward-up map. There is a wealth of information in this environment to help in navigation.

[**] Strategies within the radial grid condition tended to focus on the center landmark, moving through each "pie slice" sequentially until targets were found.

4. ENVIRONMENTAL FAMILIARIZATION

Imagine that you could visit places using a virtual environment before you actually got there. If it were a vacation spot or resort, you might like to see if you like the views, the beaches, or anything else that might catch your interest. If you were planning to drive into a city you've never been to before, you might plan a route using a conventional map and then drive your route in the virtual environment to become familiar with distances and key landmarks. If you were a soldier about to enter a hostile environment, you might use a virtual environment to rehearse a planned route and familiarize yourself with the area.

In each of these examples, the virtual environment in question is not a mere abstraction but is a representation of a real environment. In each case, we are not interested in how well the user navigates the virtual environment but rather in how they navigate the real world after exposure to the virtual environment. This section will examine several empirical studies involving the use of virtual environments for environmental familiarization. Studies of this kind are central to the topic of this chapter for two primary reasons:

1. Such investigations highlight the differences between navigation in the real world and in virtual environments, whereas studying virtual environment navigation alone does not. Studying behavioral differences between virtual and real world navigation may provide insights to help us better understand and model human navigation; virtual or real. People know how to navigate in the real world. They bring this knowledge with them into virtual environments. So, while this does not mean that virtual environment navigation must replicate real world navigation in every way, it makes sense that we need to understand how humans navigate in the real world before we can optimize methods for navigation in the virtual environment.

2. While the potential use of virtual environments to enhance training for real-world performance of many tasks has been touted from the inception of the technology, few application domains have clearly demonstrated a significant enhancement. Environmental familiarization (or mission rehearsal) is among these applications. In fact, a close inspection of the literature will lead to the conclusion that there is much confusion over whether or not virtual environments offer a significant enhancement over traditional methods of spatial knowledge acquisition. Furthermore, even if we knew they were useful, we still don't know exactly how they should be used to optimize the positive effects we want while minimizing the negative effects (e.g. reverse training) we don't want. Transfer of spatial information might be a near-term training domain that is within the reach of current technology. So, environmental familiarization may represent a microcosm of the issues involved in virtual training of any knowledge domain.

A key aspect of the studies presented in this section is that performance measures are made in the real world to evaluate how much spatial information was acquired in the virtual environment or how the virtual environment may have affected behavior. This is not to say that measurements are never made in the virtual environment, only that without real world measurements, we cannot know what affect, positive or negative, the virtual environment tool may have had on the participants.

4.1 Spatial Knowledge Transfer Studies

To ground this investigation, we compare four actual studies. There are many more studies of this type in the literature but we have selected a subset for presentation here. Each of these is a transfer study implying that there is a training phase involving a virtual environment of some type and a testing phase involving transfer to the real environment. We are also only interested here in studies about environmental familiarization rather than skill development so the virtual environment must replicate a specific real environment rather than some generic real environment such as learning to navigate in a generic city.

The purpose of this comparison is to systematically point out the similarities and differences in the studies so that we can make some statement about what is currently known about the use of virtual environments for environmental familiarization. An issue that will become very clear is that there is little consistency in the literature about what to study or how to study it. Consequently, we see a variety of experiments controlling a variety of parameters but in a way that it is difficult, if not impossible, to leverage off of what was done previously.

We will briefly discuss four experiments and conclude with a discussion of environmental familiarization and how virtual environments might be used for these tasks. The experiments are:

- Witmer, Bailey, and Knerr 1996, *Virtual Spaces and Real World Places: Transfer of Route Knowledge*
- Darken and Banker 1998, *Navigating in Natural Environments: A Virtual Environment Training Transfer Study*
- Koh, von Wiegand, Garnett, Durlach, and Shinn-Cunningham 2000, *Use of Virtual Environments for Acquiring Configurational Knowledge About Specific Real-World Spaces: Preliminary Experiment*
- Waller, Hunt, and Knapp 1998, *The Transfer of Spatial Knowledge in Virtual Environment Training*

4.1.1 A Basis for Comparison

In an attempt to make a meaningful comparison between these experiments, we will look at each in terms of a structured set of parameters. The key elements for this comparison are 1) the characteristics of the human participants, 2) the characteristics of the environment, 3) the characteristics of the tasks to be performed, 4) the characteristics of the human-computer interface, and 5) the characteristics of the experimental design to include dependent measures. Table 1 is a summary of the four experiments in terms of these elements.

EXPERIMENT	PARTICIPANTS	ENVIRONMENT	TASKS	INTERFACE	MEASURES
Darken & Banker	Domain expertise	Natural, unstructured	Maps used, 60 minute exposure, route planning and execution	Keyboard, mouse, desktop display	Deviation from route
Witmer, Bailey, & Knerr	No domain expertise	Architectural, structured	No maps used, 15 minute repeated exposure, route replication	Immersive display, buttons	Wrong turns, bearing/range estimation, time
Koh, Durlach, & von Wiegand	No domain expertise	Architectural, structured	No maps used, 10 minutes exposure, survey knowledge	Both immersive and desktop displays, joystick	Bearing/range estimation
Waller, Hunt, & Knapp	No domain expertise	Structured	No maps used, 2 or 5 minutes exposure, route replication, path integration	Both immersive and desktop displays, joystick	Time, bumps into walls

TABLE 1. A comparison of training transfer studies.

While each experiment investigated human navigation performance, the characteristics of the participant sample must be taken into account. Traditionally, important issues such as the quantity, age, and gender of the participants are considered. However, navigation is a specialized task, so we suggest that two specific differences may be quite critical. First, some of the experiments situated the navigation task within a higher-level task context. In those cases, participants' experience in the respective domain could be expected to influence both motivation to participate and task competence level. Secondly, individuals enter the experiment with a given level of innate spatial ability. Although measures of individual spatial and navigation ability may provide ambiguous results, attempts to quantify and categorize participants based upon individual ability can help to explain differences in performance.

The real-world environment that is modeled in the virtual environment influences navigation behavior. Just as individuals possess differing navigation ability, so do various environments afford different navigation experiences. Some real-world environments simply provide more navigation cues than others. Furthermore, some environments lend themselves to a higher model fidelity level than others. How closely a virtual environment matches its real counterpart is referred to as "environmental fidelity" (Waller, Hunt, & Knapp, 1998). However, do not assume that higher environmental fidelity must correlate with higher performance. There are other issues that are equally important.

With so many varieties of virtual environments and their associated interfaces, it is necessary to be more descriptive in terms of the specific differences between them. The devices and interaction styles used by the system provide differing levels of sensory stimulation to the user. How closely a virtual environment interface matches the interface to the real world (e.g. walking, driving, etc.) is referred to as “interface fidelity” (Waller, Hunt, & Knapp, 1998). Again, do not assume that higher interface fidelity equates to higher performance or training transfer. This has not yet been established and it is unclear if we will eventually find that to be the case.

Desktop virtual environments channel visual output to a computer monitor that rests on the desktop. Immersive virtual environments use a head-mounted unit or projection system to display the world to the user. Still, within both the desktop and immersive categories, the mix of input and output devices requires a more detailed description. The power of the system to deliver high-fidelity multi-modal output and monitor user input commands in real-time is a critical discriminator. The primary issue with interface fidelity in the studies we are concerned with has to do with the motion technique; specifically, what interaction method is used to control speed and direction of travel? Finally, some interfaces provide the user with special abilities and computer-generated tools that further differentiate one from another.

Experimental tasks, conditions, and standards differ across experiments. Since transfer studies consider both the training task and the testing task, we must consider both cases. The experimenter’s instructions to the participant will constrain task performance. So, two items of interest are the procedures and the dependent measures. Other items of consideration include the use of maps and exposure duration to the virtual environment.

While the experiments investigate a wide range of issues related to knowledge transfer, the central issue is what is learned and how is it applied to the real environment? The level of spatial knowledge acquired is of particular interest. A study that develops route knowledge but then tests survey knowledge may mislead the reader to believe that some other factor was the cause of low performance. Even if a system had the right interface and the right level of fidelity for a given rehearsal task, it can be used incorrectly resulting in poor performance on the transfer task in the real world.

4.1.2 The Experiments

Witmer et al, (1995) were among the first to show that a virtual environment could be useful for spatial knowledge acquisition. They compared a virtual environment to the real world in an architectural walkthrough application. They used an immersive display and their population was a random sampling without any expertise on the task. Motion was controlled by gaze-directed movement using buttons on the display.

The experimental protocol divided the session into four stages: individual assessment, study, rehearsal, and testing. During the individual assessment stage, participants responded to numerous questionnaires, some of which probed their sense of direction and navigation experience level. Next, regardless of experimental treatment condition, every participant was given fifteen minutes to study written step-by-step route directions and color photographs of landmarks. Half of the participants in each treatment condition were also provided a map of the building as a third study aid. Following the fifteen minute study stage, each participant rehearsed the route three times, either in the virtual environment or the real world depending on their group. All participants were required to identify six landmarks on the route, and researchers provided immediate correction if a participant made a wrong turn or misidentified a landmark. Finally, participants were tested in the real building. They were asked to replicate the route they had learned and to identify the six landmarks. The route replication measures included attempted wrong turns, route traversal time, route traversal distance, and misidentified landmarks. Configurational knowledge was tested by requiring participants to draw a line on a map from their known location to an unseen target.

This study effectively showed how landmark knowledge can become route knowledge but survey knowledge was not given the opportunity to develop. Survey knowledge takes time to develop by primary exposure to an environment. Exposure times for this experiment were not long enough for this to occur. Nevertheless, this study effectively provided optimism that the technology could work for this purpose – but not how well or under what conditions.

Darken and Banker (1998) studied how a virtual environment might be used as an augmentation to traditional familiarization methods. They compared performance of three groups; a map only group, a virtual environment group that also had the use of the map, and a real world group that also had the use of the map. The interface to the virtual environment was a desktop display controlled with a keyboard. The environment used was a natural region of central California with a few man-made structures but largely vegetation and rough paths (See Figure 10). They used a participant population with specialized knowledge of the task, specifically sport orienteers and experienced military land navigators.

FIGURE 10. The top image is a photograph from the testing area in the former Ft. Ord, California. The bottom image is a snapshot from virtual Ft. Ord at that exact location.

The experimental session was comprised of two phases: planning/rehearsal and testing. During the planning/rehearsal portion of the session, which lasted sixty minutes, participants studied and created a personal route from the starting point to nine successive control points. By the end of the planning phase, participants were required to draw their planned route on the map. Testing involved execution of the planned route in the real environment without the aid of the map or compass. As the participant navigated the real world course, the researcher followed, videotaping participant behavior with a head-mounted camera. A differential GPS unit worn by the participant recorded their position every two seconds. This information was used to measure the quantity of unplanned deviations from the route and the total distance traveled. In addition, the frequency of map/compass checks was recorded as a dependent measure.

This study is unique in many respects. It required the participants to plan their own routes rather than practice a given one. This serves to develop survey knowledge since alternative routes must be explored. They also attempted to introduce individual experience as a factor in addition to spatial ability. Since the task was specific to a particular domain, experience on these types of tasks should have an effect.

The results show that only intermediate participants seem to improve with the use of the virtual environment. Beginners have not yet developed enough skill at the task to be able to make use of the added information the virtual environment offers, and advanced participants are so highly proficient at map usage that the virtual environment simply does not add much information they cannot gain in other ways. On a complex environment such as this, even the hour of exposure provided does not seem to be enough. Given very short train-up times, maps still seem to be the best alternative for spatial knowledge acquisition.

Koh et al (2000) expanded previous work by specifically looking towards the development of configurational knowledge in architectural environments. In their experiment, they compared a real world group, an immersive virtual environment group, a desktop virtual environment group, and what they call a “virtual model” group that is similar to a non-interactive World-In-Miniature that is held in the hand. They used a general population sample and varied the interface to the virtual environment as described by the group (immersive or desktop). In the two virtual environment conditions, participants controlled their motion direction and speed with a joystick. The desktop group viewed a typical computer monitor. The immersive group members wore a head-mounted display. Their heads were tracked, and head rotation updated the visual scene although gaze direction was not linked to the direction of travel.

The experimental sessions were split into three phases: administration, training, and testing. During the administrative phase, the participants were informed about the bearing and distance estimation task, although the specific stations and targets were not disclosed. Participants in the three non-real world groups then were provided a period of time to familiarize themselves with the interface. Training consisted of ten minutes of free exploration under the assigned treatment conditions. The members of the real world group explored the real building and members of the virtual environment groups explored the synthetic building. Testing was specific to configurational knowledge. Participants were asked to estimate distances and bearings to unseen targets in the real world while being transported from place to place while blindfolded in a wheelchair.

The purpose of this study was to specifically focus on the development of configurational knowledge. The experimenters were not interested in landmark or route knowledge, although they do reference the hierarchical landmark, route, survey model. Furthermore, the researchers pronounced a bias that higher fidelity experiences may not necessarily lead to better transfer of knowledge, hence the use of three different fidelities in different configurations. Their results show that the virtual environment conditions do develop configurational knowledge at a comparable level to the real world.

Waller et al, (1998) used six different conditions based on the practice method. The conditions were real-world, map, no-study, and three different virtual groups -- desktop, short-duration immersive, and long-duration immersive. The environment was a maze constructed of full length curtains with targets placed at selected locations. They also used a general population sample with no specific experience in these tasks.

All participants in virtual environment conditions controlled their motion using a joystick, and the immersive groups both used a head-mounted display. The short-duration immersive group spent a total of twelve minutes in practice, while the long-duration immersive group spent a total of thirty minutes in practice over repeated trials.

The experimental protocol is comprised of four phases. The first was administrative and included proctoring of the Guilford-Zimmerman spatial abilities test. The second phase tested route knowledge by interleaving practice and testing six times. The participant would make a practice run followed immediately by a testing run. The third phase tested survey knowledge. The experimenter altered the configuration of the maze so that portions of the learned route were now blocked. The blindfolded participant then had to find a new route from one location to another. Dependent measures for both phases two and three included time to traverse the route and quantity of times the participant bumped into a wall. Finally, in the last phase, participants completed a pencil-and-paper test of their configurational knowledge.

This study showed that maps were just as effective as short durations of training but that if enough time was given, the virtual environment did prove to be more effective -- even to the point of surpassing the real world condition.

We attempt to explain all these results in the graph shown in Figure 11. Note that these curves are hypothetical since we cannot directly compare all the training transfer studies in the literature and even if we could, there are far too few data points to establish the shape of the curves. But based on our research, we believe that given a short exposure duration, maps are better than any virtual environment alternative simply because they do not overload the user with information that cannot be absorbed. However, the map is only so useful. Given enough time, the added information a virtual environment may provide will increase performance. The dark vertical bar represents the time slice that encompasses most of the studies in the literature. Very rarely are participants given enough time to develop survey knowledge in any meaningful way. We also differentiate a general virtual environment (VE) that is assumed to be only a virtual replication of the real space, from a specialized virtual environment (VE+) which might have added training features such as aerial views, transparent walls, or other features that may enhance the training effect. However, it is important to note that current research in this area is attempting to determine what those features are.

We only suggest here that we will eventually know what tools to use under what conditions in such a way that we can create virtual environment systems for environmental familiarization that compare or even surpass the real world.

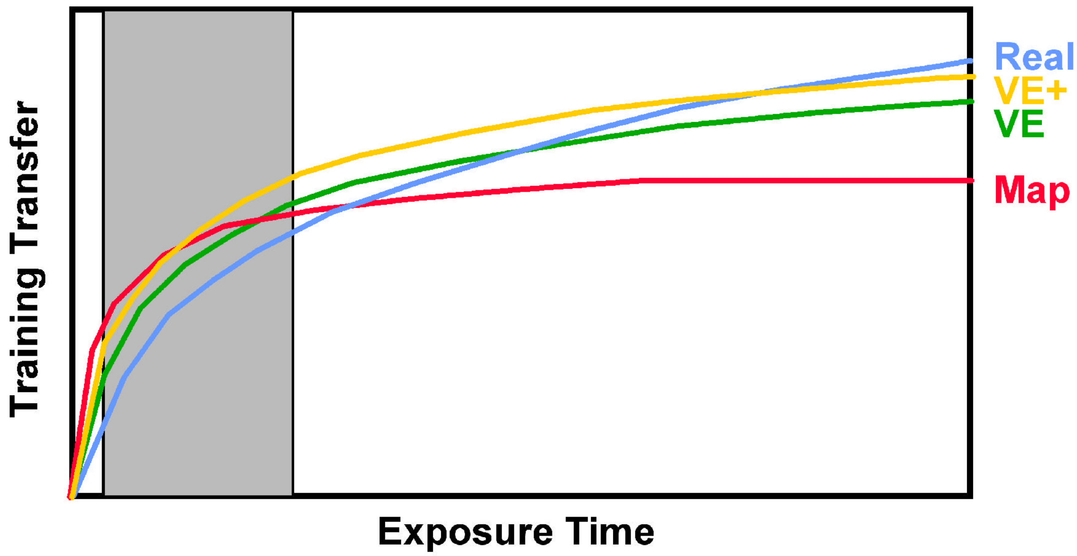

FIGURE 11. This graph shows a hypothetical picture of how spatial knowledge might be acquired over time depending on the apparatus used. Maps are best for short term events, but the real world or virtual environments with training interventions (VE+) are best over time.

5. PRINCIPLES FOR THE DESIGN OF NAVIGABLE VIRTUAL ENVIRONMENTS

This final section will serve to summarize the chapter into a series of principles we have discussed in earlier sections. This is not to be interpreted as a design "cookbook" where designers can look to see how to select navigation aids for their application. We again stress that navigation is a highly aggregate task involving people, tasks, and environments. It is not possible to develop a design solution without addressing all three elements as a whole system and not as a set of parts. What we provide here is a set of guidelines based on the literature and on our experiments as a starting point for designers to address the important issues they may face. In some cases, a technique we have discussed may fit perfectly, in others, an adaptation may be needed. The key is in providing enough spatial information so that users can execute navigation tasks as demanded by the application without overconstraining the interface, thereby eliminating exploration and discovery.

PERFORMANCE ENHANCEMENT (section 3)
TOOLS AND MEDIATORS (section 3.1)
<u>Map Usage</u> (section 3.1.1)

1. Maintain orientation but match map orientation to the task – predominantly egocentric tasks like searching should use a forward-up map while geocentric tasks such as exploration should use a north-up map.
2. Always show the user's position and view direction on the map and update dynamically.
3. The orientation problem will be more severe with a user population that includes individuals with low spatial abilities such as mental rotation – be aware of who your users are.
4. Video game play may have an effect on the selection of an appropriate map for a virtual environment. If the user population is largely a gaming community, their spatial abilities are likely to be high.
5. Use moded maps (e.g. where the use of the map precludes motion) only where appropriate. The default method should be unmoded.

<u>Landmarks</u> (section 3.1.2)

1. Allowing users to annotate the environment in some way to "personalize" the spatial cues they wish to use can be effective for complex spaces and is easily adaptable to a wide variety of navigation tasks.

2. Beware of giving the user the ability to clutter the space with excessive information. What they think will be helpful can become distracting noise.
3. Provide enough obvious landmarks (typically dependent on the context of the application and the task) so that the navigator has "reassurance" cues along a route to know that they are on the right path.
4. Make sure that landmarks are directional as well as salient – they should help provide orientation information to the navigator as well as position information.

Trails or Footprints (section 3.1.3)

1. Simply leaving a trail is marginally useful since it tends to clutter the space. Making the trail such that it disappears over time is better but can be confusing since it no longer tells the navigator that this place has been visited before, only that it hasn't been visited lately.
2. Trails can be particularly effective for exhaustive searches. It may be appropriate to use them in the context of a specific exhaustive search but then turn them off afterwards.

Directional Cues (section 3.1.4)

1. Directional cues (e.g. sun, compass) alone will not be satisfactory as a navigation aid. They should be used with other techniques since they do not provide positional information.
2. Directional cues can be effective when moded.
3. Directional cues, when used with directional landmarks are highly effective since they place landmarks in a global coordinate system.

ORGANIZATIONAL REMEDIES (section 3.2)

Implicit Sectioning (section 3.2.1)

1. Use an organizational principle wherever possible and do not violate it.
2. If the organizational principle must be violated, make it obvious where and why the violation occurred so the navigator doesn't attempt to resolve the violation into the organization principle.
3. Match landmarks to the organizational principle whenever possible. They can be used to reinforce the shape of the space.
4. Do not blindly try to use the elements of urban design in any virtual environment. They might not be appropriate. Keep in mind the key concepts: provide useful paths, observable edges, usable landmarks, and divide big, complex spaces into a number of smaller navigable spaces that are connected in some clear, understandable way.
5. Use an urban metaphor for abstract data judiciously. Make it clear where the metaphor fits and where it does not.

Explicit Sectioning (section 3.2.2)

1. Use explicit sectioning, particularly when implicit reorganization does not work for a particular environment.
2. Again, use an organizational principle and make it obvious.
3. Select a scheme for organizing your space based on what tasks people are likely to do there. If they are doing a lot of naïve searches, for example, make sure there is a way for them to easily and systematically explore the entire space without repetition.

ENVIRONMANTAL FAMILIARIZATION (section 4)

1. Do not assume that because someone can efficiently navigate a virtual environment that they can navigate the real world as well. This is simply not the case.
2. Beware of creating performance "crutches" by adding features in the virtual environment that enhance performance there but that are not available in the real world.
3. Given a short amount of time with which to familiarize someone with an environment, use a map and maybe some photos if they're available.
4. Beware of developing orientation-specific survey knowledge if only maps are used. Given enough time, a virtual environment can be used to develop orientation-independent spatial knowledge, but it takes time.
5. Be careful when deciding how to use a virtual environment for environmental familiarization because it is extremely difficult to compare studies in the literature. Make a decision based on the whole problem – the people you are training, the tasks they are doing, and the environment they are navigating in.

6. REFERENCES

Aretz, A. J., & Wickens, C. D. (1992). The Mental Rotation of Map Displays. Human Performance, 5(4), 303-328.

Boer, L. C. (1991). Mental Rotation in Perspective Problems. Acta Psychologica, 76, 1-9.

Bowman, D., Koller, D., & Hodges, L. (1997). Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques. Paper presented at the Virtual Reality Annual Internation Symposium (VRAIS), Albuquerque, NM.

Chase, W. G. (1983). Spatial Representations of Taxi Drivers. In D. R. Rogers & J. A. Sloboda (Eds.), Acuisition of Symbolic Skills, . New York: Plenum.

Chen, J. L., & Stanney, K. M. (2000). A Theoretical Model of Wayfinding in Virtual Environments: Proposed Strategies for Navigational Aiding. PRESENCE: Teleoperators and Virtual Environments, 8(6), 671-685.

Colle, H. A., & Reid, G. B. (1998). The Room Effect: Metric Spatial Knowledge of Local and Separated Regions. PRESENCE: Teleoperators and Virtual Environments, 7(2).

Darken, R. P. (1996). Wayfinding in Large-Scale Virtual Worlds. Unpublished Doctoral dissertation, The George Washington University.

Darken, R. P., & Banker, W. P. (1998). Navigating in Natural Environments: A Virtual Environment Training Transfer Study. Paper presented at the IEEE Virtual Reality Annual International Symposium, Atlanta, GA.

Darken, R. P., & Cevik, H. (1999). Map Usage in Virtual Environments: Orientation Issues. Paper presented at the IEEE Virtual Reality 99, Houston, TX.

Darken, R. P., Cockayne, W. R., & Carmein, D. (1997). The Omni-Directional Treadmill: A Locomotion Device for Virtual Worlds. Paper presented at the ACM UIST '97, Banff, Canada.

Darken, R. P., & Sibert, J. L. (1993). A Toolset for Navigation in Virtual Environments. Paper presented at the ACM Symposium on User Interface Software and Technology, Atlanta, Ga.

Darken, R. P., & Sibert, J. L. (1996). Wayfinding Strategies and Behaviors in Large Virtual Worlds. ACM SIGCHI 96, 142-149.

Donelson, W. C. (1978). Spatial Management of Information. Proceedings of ACM SIGGRAPH '78, 203-209

Downs, R. M., & Stea, D. (1977). Maps in Minds: Reflections on Cognitive Mapping. New York: Harper & Row.

Durlach, N., & Mavor, A. (Eds.). (1995). Virtual Reality: Scientific and Technological Challenges. Washington, D.C.: National Academy Press.

Furnas, G. W. (1986). Generalized Fisheye Views. Proceedings of ACM SIGCHI '86, 16-23.

Goldin, S. E., & Thorndyke, P. W. (1982). Simulating Navigation for Spatial Knowledge Acquisition. Human Factors, 24(4), 457-471.

Ingram, R., & Benford, S. (1995). Legibility Enhancement for Information Visualisation. Paper presented at the Visualization 1995, Atlanta, GA.

Jul, S., & Furnas, G. W. (1997). Navigation in Electronic Worlds: A CHI 97 Workshop. SIGCHI Bulletin, 29(4), 44-49.

Koh, G., von Wiegand, T., Garnett, R., Durlach, N., & Shinn-Cunningham, B., (2000). Use of Virtual Environments for Acquiring Configurational Knowledge About Specific Real-World Spaces: Preliminary Experiment. PRESENCE: Teleoperators and Virtual Environments 8(6), 632-656.

Levine, M., Marchon, I., & Hanley, G. (1984). The Placement and Misplacement of You-Are-Here Maps. Environment and Behavior, 16(2), 139-157.

Lynch, K. (1960). The Image of the City. Cambridge: MIT Press.

McGee, M. G. (1979). Human Spatial Abilities: Psychometric Studies and Environmental, Genetic, Hormonal, and Neurological Influences. Psychological Bulletin, 86(5), 889-918.

McNamara, T. P., Ratcliff, R., & McKoon, G. (1984). The Mental Representation of Knowledge Acquired from Maps. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10(4), 723-732.

Neisser, U. (1976). Cognition and Reality: Principles and Implications of Cognitive Psychology. New York: W.H. Freeman and Company.

Passini, R. (1984). Wayfinding in Architecture. New York: Van Nostrand Reinhold Company Inc.

Péruch, P., Pailhous, J., & Deutsch, C. (1986). How Do We Locate Ourselves on a Map: A Method for Analyzing Self-Location Processes. Acta Psychologica, 61, 71-88.

Presson, C. C., DeLange, N., & Hazelrigg, M. D. (1989). Orientation Specificity in Spatial Memory: What Makes a Path Different From a Map of the Path? Journal of Experimental Psychology: Learning, Memory, and Cognition, 15(5), 887-897.

Presson, C. K., & Hazelrigg, M. D. (1984). Building Spatial Representations Through Primary and Secondary Learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10(4), 716-722.

Rieser, J. J. (1989). Access to Knowledge of Spatial Structure at Novel Points of Observation. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15(6), 1157-1165.

Rossano, M. J., & Warren, D. H. (1989). Misaligned Maps Lead to Predictable Errors. Perception, 18, 215-229.

Siegel, A. W., & White, S. H. (1975). The Development of Spatial Representations of Large-Scale Environments. In H. Reese (Ed.), Advances in Child Development and Behavior, (Vol. 10,). New York: Academic Press.

Spence, R. (1998). A Framework for Navigation (Technical report 98/2). London: Imperial College of Science, Technology, & Medicine.

Stevens, A., & Coupe, P. (1978). Distortions in Judged Spatial Relations. Cognitive Psychology, *10*, 422-437.

Stoakley, R., Conway, M. J., & Pausch, R. (1995). Virtual Reality on a WIM: Interactive Worlds in Miniature. Paper presented at the Proceedings of ACM SIGCHI 95, Denver, CO.

Suchman, L. A. (1987). Plans and Situated Actions: The Problem of Human Machine Communication. Cambridge: Cambridge University Press.

Templeman, J., Denbrook, P. S., & Sibert, L. E. (2000). Virtual Locomotion: Walking-In-Place Through Virtual Environments. PRESENCE: Teleoperators and Virtual Environments, 8(6), 598-617.

Thorndyke, P. W., & Goldin, S. E. (1983). Spatial Learning and Reasoning Skill. In H. L. Pick & L. P. Acredolo (Eds.), Spatial Orientation: Theory, Research, and Application, (pp. 195-217). New York: Plenum Press.

Thorndyke, P. W., & Stasz, C. (1980). Individual Differences in Procedures for Knowledge Acquisition from Maps. Cognitive Psychology, 12, 137-175.

Tolman, E. C. (1948). Cognitive Maps in Rats and Men. Psychological Review, 55(4), 189-208.

Waller, D., Hunt, E., & Knapp, D. (1998). The Transfer of Spatial Knowledge in Virtual Environment Training. Presence: Teleoperators and Virtual Environments 129-143.

Witmer, B. G., Bailey, J. H., & Knerr, B. W. (1995). Training Dismounted Soldiers in Virtual Environments: Route Learning and Transfer (Technical Report 1022): U.S. Army Research Institute for the Behavioral and Social Sciences.

.