

Map Usage in Virtual Environments: Orientation Issues

Rudolph P. Darken and Helsin Cevik
Department of Computer Science
Naval Postgraduate School
Monterey, California 93943-5118
+1 831 656 4072
darken | hcevik@cs.nps.navy.mil

Abstract

Navigation tasks in large virtual environments often call for the use of a virtual map. However, all maps are not alike. Performance on navigation tasks in general has been shown to vary depending on the orientation of the map with respect to the user's frame of reference. This paper reports the results of an experiment investigating orientation issues of virtual maps for use during navigation tasks. Participants were given a virtual map in either a north-up or forward-up configuration. Performance on search tasks was measured in terms of search time and errors. Results indicate that targeted search tasks (tasks requiring only the egocentric reference frame) are best served by a forward-up alignment while primed and naive search tasks (tasks requiring information from the world reference frame) prefer a north-up alignment. Both types of maps are affected by the ability of the user to perform mental rotations.

1. Introduction

Given the importance of navigation to general task performance in large-scale virtual environments (VEs), maps have often been used as a solution to wayfinding problems. However, the actual benefits and limitations of map use in this new medium are largely unknown. It seems that in most cases, any map is better than no map, especially for very large VEs that are sparsely populated or otherwise generally difficult to navigate. But if a map is to be used, how should it be presented to be most effective?

There are underlying principles, founded primarily in cognitive psychology, geography, and cartography that can help us determine how maps should be presented to achieve optimal performance on navigation tasks. However, maps in VEs differ significantly from maps in the real world. Because the viewpoint location is known at all

times, virtual world maps can dynamically show the position of the viewpoint (a.k.a. the You-are-here or YAH position) during navigation rather than only at static locations typical of shopping mall maps. The same can be said of dynamic objects in the environment. This trivializes the transformation of positions in the egocentric reference frame (ERF) to the world reference frame (WRF). It does not, however, trivialize the rotations necessary to align the ERF with the WRF.

This paper presents the results of an experiment to determine how map orientation affects different navigation tasks in large VEs and the implications of this phenomenon on map usage in VEs in general.

2. Factors Affecting Map Orientation

The problem of determining optimal map orientation is complicated by conflicting task demands. Exploration does not necessarily require active navigation. If the goal is to extract spatial knowledge from the environment (i.e. develop a mental representation or cognitive map) this can be achieved to some extent by map study alone. Searching a space, however, assumes some target object or location and *does* require active navigation. Searching tasks use the cognitive map along with search strategies or heuristics to constrain the search wherever possible. In its tendency to develop a cognitive map, exploration tasks improve performance on subsequent searching tasks. The reverse is also true. Repetitive searching tasks in the same environment will tend to develop a cognitive map just the same as an exploration task, possibly better. While searching tasks are most prevalent in games such as DOOM™* and its derivatives, there is also a strong demand for configuration knowledge†. Players who learn optimal paths and alternate paths through these environments have an advantage over

*. © Id Software, Inc.

†. Topological, map-like information

those who are unfamiliar with the space. Also, the space itself clearly has an effect on performance. Large, sparse spaces have fewer cues with which to perform a triangulation and therefore are difficult to navigate. On the other extreme, dense, overpopulated spaces may have too much complexity, also resulting in decreased navigation performance.

Map use during navigation is fundamentally different from map use explicitly for spatial knowledge extraction. During navigation, a transformation from the ERF to the WRF is *required*. It is here where we see people turning maps in different directions while they are driving or walking in an unfamiliar area. This strategy is not used while looking over a map to plan a trip, for example. Concurrent navigation with map use must simplify the ERF to WRF transformation to be useful in practice.

There is also an issue involving navigation tasks performed on the map rather than on the environment itself. When in map mode, DOOM™ players can move themselves in the environment by manipulating the YAH marker on the map (see figure 1). The Worlds-In-Miniature (WIM) metaphor takes this further by allowing navigation on the WIM (WRF) while simultaneously updating the ERF perspective [1]. No data exists as to the effectiveness of either of these maps.



Figure 1: The north-up DOOM™ map with YAH marker.

3. Theoretical Background

Map orientation has been an important topic of research in the aviation community for many years due to their concern with pilots' ability to successfully merge a map display (WRF) with the "out-the-window" view (ERF). The two primary categories of map orientation are *north-up* which, as the name implies, always has north at the top, and *forward-up* (or *track-up*) which is always rotated so as to be congruent with the forward direction.

There are instances of both of these types of maps both in the video game industry and VEs. The DOOM™ map (see

figure 1) is a north-up map but can be translated horizontally, vertically or zoomed by the user. It cannot be rotated. Many other games use the forward-up orientation. NPSNET uses a forward-up targeting display in its center with a compass in the lower left corner showing field of view (see figure 2) [2]. The targeting display shows targets in their relative locations in the ERF.




Figure 2: The forward-up NPSNET targeting map.

The primary cognitive factor related to map orientation is that of *mental rotation*. When the ERF and the WRF are misaligned, a mental rotation must take place to bring them into congruency. Aretz and others have suggested that there is a linear relationship between the magnitude of misalignment and the time it takes to perform the rotation [3]. However, in studying this phenomenon, most experiments have failed to view navigation as a process rather than as a series of discrete perspective transformations [3-5]. Maps used for navigation *during* navigation tasks maintain context and consistency that snapshot images do not preserve.

Mental rotation is also a factor when navigating from a cognitive map. As a mental representation of a space is developed, it tends to be either orientation-independent or orientation-dependent, usually based on the source of information. It has been shown that spatial knowledge acquired from a north-up map tends to be orientation-specific while spatial knowledge acquired from first-hand active navigation tends to be orientation-independent [6]. For example, after studying a north-up map of a city, when entering that city from the north, the cognitive map will be 180° out of alignment causing a mental rotation to occur. Obviously, orientation-independency is preferable but often takes significant time to develop.

The goal of this research is to identify principles for map presentation that will enhance navigation performance in a general way. Ideally, performance would be comparable to that of an experienced navigator of the environment (i.e. someone with a well-developed orientation-independent cognitive map).

4. Approach

This study was intended to investigate effects of map orientation on navigation performance toward a set of principles to assist virtual world builders with map design issues. However, we needed to widen the scope of our experiment in order to better understand the problem. It is clear that, beyond the orientation of the map, other factors come into play.

The individual differences of users, particularly with respect to mental rotation, are a critical concern. We are also interested in whether or not the orientation of the map might influence strategy. Another issue is the type of space — specifically its spatial characteristics such as relative size, object density, etc. Lastly, we are interested in how performance of different navigation tasks are affected by all of these factors. Different tasks require different types of information and consequently, should be affected by map orientation and individual differences.

5. Method

Our experiment examined users of an urban and an open ocean VE executing a set of navigation tasks with each of two maps with different orientation schemas.

Navigation tasks are coarsely defined to be either searching or exploration tasks [7,8]. For this study, the following navigation tasks were used:

- *Targeted search*: A searching task in which the target in question is shown on the virtual map.
- *Primed search*: A searching task in which the location of the target is known, but the target does not appear on the virtual map. The search is presumed to be non-exhaustive.
- *Naïve search*: A search task in which there is no a priori knowledge of the whereabouts of the target in question and the target is not shown on the map. A naïve search implies that an exhaustive search must be performed. An optimal exhaustive search requires that the navigator traverse the entire space once (in the worst case).
- *Exploration*: A wayfinding task in which there is no specific target.

Exploration was not explicitly examined in this study but it is clearly intertwined with searching tasks to some degree. In addressing the needs of all of these categories, virtual maps must support both exhaustive and non-exhaustive searches and must facilitate acquisition of configuration knowledge.

The performance of each participant was observed and measured on each navigation task in both the urban and

open ocean VEs. While each participant received a condition in both types of environment, they only received one type of map treatment. The order of their treatments was predetermined at random. The two map treatments in the study were:

- *Forward-up Treatment*: The orientation of the map remains static with respect to the world. The forward direction in the world is always at the top. The YAH marker is represented as a sphere. See figure 3.
- *North-up Treatment*: The orientation of the map remains static with respect to the participant. The top or “north” is always at the top. The YAH marker is represented as a cone with origin at the viewpoint. See figure 4.

We considered a third map condition, a *manually-oriented treatment*, where the map is not automatically oriented in any fashion but rather is rotated via a tracker held in the user’s hand. This was of interest because it most closely resembles the use of paper maps in the real world. We wondered whether users would attempt to replicate the forward-up map or the north-up map. The answer seems to be neither. The task of searching for targets, when combined with the task of keeping the map aligned properly, seems to be more than the average person can handle. They quickly lose track of how the map should be oriented and consequently, they stop using it. It was determined from the results of our earlier pilot study that this type of map was far inferior in terms of performance to the other two types of maps on all categories of search tasks. We, therefore, did not continue its investigation in this study.

In order to determine the effects of differences in spatial ability on navigation performance, participants were given the Guilford-Zimmerman spatial visualization (SV) and spatial orientation (SO) standardized tests prior to beginning the VE trials.

Before beginning the VE portion of the experiment, each participant was given a period of time on a separate practice VE to become familiar with the Fakespace Inc. PUSH™ device and the movement mechanism inherent to that device. During trials, the paths each participant traversed through each assigned navigation task were sampled (approximately once per second) for later analysis and task completion times were taken. During task execution, participants were asked to “think aloud” [9] as a method of knowledge elicitation specifically aimed at understanding search strategies. In addition to task completion times, we analyzed errors, defined as “wrong turns” where participants turned away from the target rather than towards it, presumably due to mental rotation errors.

The navigation tasks performed for all treatments required the participant to execute three targeted searches, followed by three primed searches, followed by one naïve




Figure 3: The Forward-Up treatment in the urban environment.

search. For the targeted searches, a colored target (red, green, and blue) appeared on the map, one at a time. Then, the participant was instructed to return to each of the three previous targets but no visual feedback was given on the map. These are the primed searches. Finally, the participant was asked to locate a target not seen before and with no visual feedback on the map. This is the naïve search.

5.1. Design

Thirty participants (4F / 26M) took part in the study. They were divided into two groups of fifteen at random with the constraint that each group be balanced with respect to spatial ability scores. All participants had a technical background and were between the ages of 25 and 36.

5.2. Stimuli and Apparatus

There were two VEs used for this study: an urban VE (UVE) and an open ocean VE (OOVE). For the UVE, we used a modified version of Performer* town. The OOVE was manually constructed using a geometric modeling tool. The OOVE contains large areas of open sea with four land masses. The land masses are shaded by elevation and the ocean surface is textured. The land masses were shaped and scaled to be distinct from one another. The targets (colored spheres) were manually placed in the worlds. These environments were selected due to their spatial characteristics. We were interested in whether or not objects in a space (or lack thereof) would have an effect on performance or strategy.

The viewpoint could not be moved vertically in either environment and was fixed at three meters in elevation. Movement was also constrained horizontally so that when

*. Silicon Graphics, Inc.




Figure 4: The North-Up treatment in the urban environment.

an edge was encountered, it could not be used as a navigation aid.

The virtual maps used in both treatments were identical to their corresponding environment. A YAH marker was moved along the map surface to identify the viewpoint and to establish location in the actual VE. The YAH marker was a white sphere in the forward-up map treatment and a red cone approximating the view volume for the north-up map treatment. This difference was due to the fact that the forward-up configuration preserves relative direction information while the north-up configuration does not. However, this fact in and of itself, does not account for any obvious benefit as long as the YAH marker in the north-up treatment shows view direction.

The virtual maps were presented tilted by 45° and above the actual VE. The maps were unmoded, meaning that participants can view both the map and the forward view simultaneously. The maps were not interactive and could not be directly manipulated by the participant.

The treatments were implemented on a Silicon Graphics Onyx™ Infinite Reality™ graphics workstation. A Fakespace Inc. PUSH™ display and tracker was used for head tracking and visual display. The PUSH™ is a full color, high-resolution device that provides full six-degree of freedom movement in the VE. The display is held to the eyes with two hands. The position and orientation of the head are tracked through three mechanical joints. Motion is controlled via an acceleration metaphor by which the participant accelerates forward in the VE by pushing the display forward or decelerates by pulling the display backwards. Movement can be forward or backward but is always in the direction of view. The participant may stop at any time by releasing the display.

6. Results and Discussion

The images in the following sections are examples from the experimental data, chosen as typical of their group. In the urban environment, the search trail is indicated by a stream of white dots. In the ocean environment, it is indicated by a thick line.

6.1. Effects of Map Orientation

With a 0.1 level of significance, a two-sample t-Test ($t=1.59$; $P < 0.1$) concludes that the forward-up map orientation is superior to the north-up map orientation for targeted search tasks in the UVE with regard to time and errors (see figure 5). Targeted searches are viewed as purely in the




Figure 5: Time and errors in targeted searches in the forward-up and north-up treatments in the UVE.

ERF, and consequently, only ERF information is required. This is best supplied by the forward-up map because it requires no mental rotation. As seen in figure 6, the forward-up track (A) is far more direct than the north-up track (B). The north-up map tends to be susceptible to “wrong turns” where a user goes left when they mean to go right due to an error in transforming the ERF to the WRF on the map.



Figure 6: Targeted searches in the forward-up (A) and north-up (B) treatments.

This result carries over from dense environments to sparse environments where, with a 0.01 level of significance, a Wilcoxon Rank Sum nonparametric test ($P < 0.01$) concludes that the forward-up map orientation is better than the north-up map orientation for targeted search tasks in the OOVE (see figure 7).




Figure 7: Time and errors in targeted searches in the forward-up and north-up treatments in the OOVE.

With a 0.1 level of significance, we cannot conclude that one map orientation is better than the other for primed search tasks. However, further analysis shows that participants in the forward-up map group with high spatial abilities perform better than participants in the north-up group, also with high spatial abilities. Conversely, participants with low spatial abilities in the forward-up map group perform worse than participants in the north-up map group, also with low spatial abilities. This effect resulted in a balance between the scores of both groups. What seems to happen is that high spatial individuals treat primed searches like targeted searches because they have few difficulties in transforming the ERF to the WRF and vice versa. Individuals with low spatial abilities, however, have greater difficulties in transforming the ERF to the WRF, and consequently often treat the primed searches like naïve searches, e.g. their strategy involves performing an exhaustive search because they often cannot identify the location on the map where they need to go. This is evident in their search tracks where we often see them crossing their path and re-searching areas under the forward-up condition (see figure 8A), versus the north-up condition, where we see relatively coherent strategies and better performance (see figure 8B).

The naïve search seems to show a similar effect, but not at significant levels. The forward-up map (see figure 9A) illustrates an unorganized exhaustive search, typical of this treatment. Since the map seems to “turn” in front of the user (actually it is the user turning about the map), it is difficult to develop a reasonable strategy to cover the entire environment in an efficient manner. The north-up map, however, (see figure 9B) illustrates a relatively efficient




Figure 8: Primed searches of low spatial participants in forward-up (A) and north-up (B) treatments.

search. This participant happens to have been fairly lucky in finding the target after searching only two blocks of the town, but it is safe to say that had the target been somewhere else, it would have been located in short order based on the “back and forth”^{*} method used here.



Figure 9: A naïve search in the forward-up (A) and north-up (B) treatments.

These same results are also true of the OOVE. With a 0.1 level of significance, we cannot conclude that one map orientation is better than the other for primed search tasks. However, the same difference with respect to high and low spatial individuals applies here. We noticed that the forward-up map (see figure 10A) tended to exhibit slightly poorer task execution than the north-up map (see figure 10B). In this particular case, the forward-up map shows an error around the target in the lower left corner where the north-up map is direct and accurate.

6.2. Effects of Individual Differences

Our results indicate a difference in performance in the north-up map condition based on spatial visualization (SV) ability. For the UVE, with a 0.05 level of significance, we conclude that participants’ spatial ability has an effect on performance for all navigational tasks (see figure 11). As seen in figure 12, the low spatial individual (A) has great difficulty transforming the ERF to the WRF and conse-




Figure 10: A primed search in the forward-up (A) and north-up (B) treatments in the open ocean environment.




Figure 11: Simple regression of SV and avg. time for targeted and primed searches in north-up condition.

quently makes many errors in locating the target. The high spatial individual (B) exhibits none of this behavior. However, for the OOVE, this effect is not as great. With a 0.25 level of significance, it cannot be definitively concluded that spatial ability has an effect on navigational performance. Nevertheless, we believe that a sample population of other than all technical individuals would have proven significant in this case.



Figure 12: A primed search in the north-up treatments of a low SV (A) and a high SV (B) participant.

^{*}. Referred to as the “Lawnmower method” in [7,8].

Our results also indicate a difference in performance in the forward-up condition in both environments based on SV and SO scores (see figure 13). This is a stronger effect,




Figure 13: Multiple regression of SV/SO and avg. time for targeted and primed searches in forward-up condition.

being consistent across both types of environment. As stated earlier, low spatial participants tend to approach both the primed and naïve searches as exhaustive searches, thus greatly lowering performance. Exhaustive searches are best served by a WRF perspective (e.g. the north-up map) because more than simple ERF directional information is required. This would imply that for the north-up condition, low spatial participants struggled with the first three targeted searches, but then improved over the rest of the trial. However, low spatial participants in the forward-up condition would have done well on the first three targeted searches and then struggled with the remaining four searches. This is in fact why the effect is stronger in our experiment. Had we divided the tasks evenly between ERF and WRF tasks, we suspect the effect would have been equally strong across map orientation conditions and environment types.

6.3. Effects of Virtual Environment Type

We had expected from the onset that navigating the OOVE would be significantly more difficult than navigating the UVE. We thought that having cues to navigate by must be better than having none, which is almost the case in our sparse ocean environment. Surely, having many cues by which to triangulate and perform mental rotations must be easier than not having those cues at all. However, the abundance of cues in the UVE seems to be exactly what may have caused problems on some search tasks. In particular, the Performer™ town environment is somewhat symmetrical. Consequently, many participants had difficulty remembering where targets had been found in order to execute primed searches in the forward-up treatment. Also,

even in cases where these cues were a benefit, these same cues are a cause of visual obstruction and they inhibited movement. Navigation performance, error rates, post-trial comments of participants, and direct observations indicate that navigating in the OOVE may actually be easier in some cases regardless of map orientation. The OOVE is very simple. Even participants with less spatial ability completed the tasks with about the same level of proficiency in the OOVE as compared to the UVE.

6.4. Effects of Game Play

Of our four participants who were avid video game players, all four performed better than average using both types of map in either environment.

There is an interesting paradox in user performance versus user preference. Many participants preferred the north-up condition over the forward-up condition although their performance did not necessarily reflect their partiality. In post-test debriefing, participants indicated that their preference was largely due to familiarity with north-up map use in games such as DOOM™. Variance on ERF task performance and WRF task performance between conditions was also lower for these participants suggesting that they are better able to transform reference frames than those who are inexperienced with this cognitive task. Although it has long been affirmed that spatial ability is not trainable, video game play seems to have some effect on it. This will be an important topic for further study considering the vast similarities between VEs and video games.

7. Conclusions

The three basic principles of virtual map presentation identified in this experiment are:

1. For ERF tasks such as targeted searches, a forward-up map is preferable to a north-up map.
2. For WRF tasks such as primed or naïve searches, a north-up map is preferable to a forward-up map.
3. Under almost every possible condition, individuals with high spatial abilities will be able to use either type of map better than individuals with low spatial abilities.

Furthermore, we found that these principles apply across types of environment 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 task versus the spatial ability of their users.

Maps that adhere to these principles and that dynamically show the viewpoint position on the map simplify ERF

to WRF perspective transformations that are required for map use.

The results of this study suggest that perspective transformation (ERF to WRF and vice versa) may be a partially trainable skill. While it is not suggested that individuals with low spatial ability can be trained to be comparable to those with high spatial ability, they may be able to raise their general level of performance with repeated exposure to perspective transformation tasks. This supposition shouldn't be taken too far based on this data alone, however. It is imperative that this experiment be followed with another studying a larger population sample, not only in sample size but also with varying spatial abilities. It may be found that individuals with particularly low spatial abilities such as mental rotation and triangulation will not show the benefits of repeated exposure to these maps and environments that individuals with high spatial ability will show.

There seemed to be several problems in using the forward-up maps with regard to losing track of where "north" was. It may be possible to significantly raise performance in the forward-up group simply by marking one side of the map in some obvious way. It has been previously shown that a global direction cue, such as a sun, can have a pronounced effect on performance and strategy [10]. This may also be true as applied to maps.

Future research in this area must address the issue of game playing experience by participants. The fact that some participants have played possibly hundreds of hours on games that use maps in different configurations with YAH indicators must have an effect on task performance, and consequently, experimental data. How this exposure or training effect interacts with natural spatial abilities is unknown at this time but warrants further study.

For tasks requiring spatial knowledge acquisition of a specific space, prolonged exposure to a VE with (and without) a virtual map display will allow us to determine if an orientation-independent cognitive map can be constructed as effectively in a virtual environment as it is in the real world allowing people to familiarize themselves with places they have never before visited.

It is our hope that a greater understanding of human navigation in VEs will lead to better, more usable VE applications and training tools in the near future.

8. Acknowledgements

This research is supported by the Office of Naval Research, Cognitive and Neural Sciences Division. Special thanks go to the NPSNET Research Group at the Naval Postgraduate School for their valued assistance and cooperation and also to our experimental participants for their time and effort.

9. References

1. Stoakley, R., M.J. Conway, and R. Pausch. *Virtual Reality on a WIM: Interactive Worlds in Miniature*. in *Proceedings of ACM SIGCHI 95*. 1995. Denver, CO: ACM Press. p. 265-272.
2. Zyda, M.J., et al., *NPSNET: Constructing a 3D Virtual World*. 1992 Symposium on Interactive 3D Graphics, 1992: p. 147-156.
3. Aretz, A.J. and C.D. Wickens, *The Mental Rotation of Map Displays*. *Human Performance*, 1992. **5**(4): p. 303-328.
4. Levine, M., *You-Are-Here Maps: Psychological Considerations*. *Environment and Behavior*, 1982. **14**(2): p. 221-237.
5. Péruch, P., J. Pailhous, and C. Deutsch, *How Do We Locate Ourselves on a Map: A Method for Analyzing Self-Location Processes*. *Acta Psychologica*, 1986. **61**: p. 71-88.
6. Thorndyke, P.W. and B. Hayes-Roth, *Differences in Spatial Knowledge Acquired from Maps and Navigation*. *Cognitive Psychology*, 1982. **14**: p. 560-589.
7. Darken, R.P. and J.L. Sibert, *Wayfinding Strategies and Behaviors in Large Virtual Worlds*. *ACM SIGCHI 96*, 1996: p. 142-149.
8. Darken, R.P. and J.L. Sibert, *Navigating in Large Virtual Worlds*. *International Journal of Human-Computer Interaction*, 1996: p. 49-72.
9. Ericsson, K.A. and H.A. Simon, *Protocol Analysis: Verbal Reports as Data*. 1993, Cambridge, MA: MIT Press.
10. Darken, R.P. and J.L. Sibert. *A Toolset for Navigation in Virtual Environments*. in *User Interface Software and Technology*. 1993. Atlanta, Ga.: ACM Press.