# Wayfinding in a Virtual Environment

| Kellogg Booth | Brian Fisher | Steven Page | Colin Ware | Sherri Widen |
|---|---|---|---|---|
| Computer Science | Computer Science | Computer Science | Computer Science | Psychology |
| U. British Columbia | U. British Columbia | U. New Brunswick | U. New Brunswick | U. British Columbia |
| ksbooth@cs.ubc.ca | fisher@cs.ubc.ca | spage@cs.ubc.ca | cware@unb.ca | ksbooth@cs.ubc.ca |

## Abstract

The goal of this study was to investigate spatial learning in a large-scale virtual environment. Our hypothesis was that participants who had actively piloted a virtual vehicle through the terrain (Drivers) would outperform passive learners (Passengers) in both route memory (drive along the path without a guide) and global terrain knowledge (drive between novel pairs of stations in the same landscape). Our results showed the opposite effect, suggesting that attentional limitations may have played a larger role in learning a virtual environment than active versus passive learning mode.

## Introduction

The goal of this study was to investigate spatial learning in a computer-rendered 3D virtual world. Large-scale virtual environments require users to gain familiarity over time with the relative locations of landmarks in order to reliably move from their current location to a desired one -- an ability that we call "wayfinding". Our working hypothesis was that participants who had actively piloted a virtual vehicle through the terrain (Drivers) would outperform passive learners (Passengers) in both route memory and global terrain knowledge.

After a brief familiarization with the control apparatus and display, participants moved along ten different paths between ten clearly marked "stations". Half of the participants were "Passengers" whose observation point automatically followed a computer-driven "pilot vehicle". The other half ("Drivers") were asked to use the control device to actively follow the pilot vehicle along the same sequence of paths.

The primary goal of this study was to examine the impact of passive versus active learning on two spatial tasks: in the *Repeating* task the route memory of both groups was tested by asking them to actively drive along the same ten paths without a guide. In a subsequent *Wayfinding* task, both groups of participants were asked to find their way between novel pairs of stations as a test of their memory of the overall structure of the terrain (i.e. their mental map) .

Researchers have hypothesized that spatial learning proceeds along a sequence of stages. First observers learn *landmarks*, later *routes* through the terrain, and finally achieve an overall - *survey* knowledge or mental map of the region (Siegel & White, 1975). In this method, participants become familiar with landmarks in the world, join sequences of familiar objects together to learn routes within the world, and finally gain an overall understanding of the entire region. This is typically thought of as forming a "mental map" of the region. The level of correspondence between an observer's mental map and the physical space is the subject of some debate. Some have argued that a mental representation of a space mirrors the spatial characteristics of the environment in the same direct way as distances and directions in a true map relate to those in the world (Kosslyn, 1981) . Others have postulated a more limited mental representation in which metric distances are relatively unimportant, and wayfinding occurs by means of specialized heuristics guided by perceptual input (Pylyshyn, 1981).

Regardless of the outcome of this debate, it is clear that the task of learning one's way around a particular environment is qualitatively different from learning a fact (semantic memory) or remembering an event in the past (episodic memory). The key difference relates to the way that memories are formed over time and are mediated by action. In both the physical world and a large virtual world, the information that can be taken in at a glance is usually insufficient to establish either the observer's absolute location in space or the angular direction and distance from a particular landmark. Instead, a sequence of views of the environment is created as the observer explores the space. These views and the actions that produced them must somehow be combined to create an effective representation and mechanism for navigating in that environment.

Most virtual environments differ from familiar physical environments in a number of characteristics that may differentially affect the process of forming mental representations and mechanisms for wayfinding: Virtual environments are typically sparse, and the relative shortage of local landmarks (such as buildings and signposts) and more global visual contexts (such as a backdrop of mountains) may impede learning.

In the physical world, sparse environments such as these are encountered in rare circumstances, e.g. navigation on the open sea without instrumentation. Wayfinding in these unusual situations often requires new methods that are specific to that environment. For example, Micronesian navigators are adept at finding their way across vast stretches of open ocean in small sailing canoes. They accomplish this difficult feat without maps or compass, by using a combination of subtle perceptual cues such as water colour and bird behaviour, and a set of learned mental "tricks". These unusual mental operations make a number of assumptions about the environment that are known to be untrue, such as postulating the presence of place-holder islands that are known not to exist, and a moving ocean underneath the static boat rather than the actual converse.

Such stratagems create an alternative mental model that does not attempt to construct a representation of the limited structure of the ocean environment. Rather than mentally mapping the space, these navigators utilize a performance model that has been developed specifically for open-ocean navigation (Hutchins 1996). Given the lack of environmental cues and navigation tools, these models have proven superior to a more realistic mental map of the space in the most direct way possible -- by the success of these ocean voyages and the survival of the voyagers. By analogy, it is possible that the most effective methods of wayfinding in sparse virtual environments will prove to be methods that are not extensions of our ways of navigating the physical world, but rather are specific to these unfamiliar spaces.

On the positive side, the addition of structure to the virtual environment may increase the utility of direct sensory input. Regularities in the structure of artificial environments may allow observers to make simplifying assumptions about the nature of the world, in the same way that the layout of streets and buildings makes wayfinding in a city easier than in a wilderness area for those who are familiar with the rules and regularities of cities.

Studies of spatial learning in cities have led to recommendations for the design of those spaces (Lynch, Banerjee. and Southworth, 1995). Some of these ideas have been taken up by designers of virtual environments (see Vinson 1999), although their effectiveness has not been established. The nature of the regularities that may be exploited also remains to be determined.

Some studies have suggested that users have a good deal of flexibility about metric characteristics of Webspaces and their representations, so violations of those aspects of the environment are well tolerated. Violations of other spatial constraints such as Gestalt item groupings and individual item constancy are hypothesized to be more likely to disrupt spatial understanding. This has led to the use of spatial transformations in Web visualization environments (Fisher, Agelidis, Dill, Tan, Collaud, and Jones 1997) that violate some of the constraints of physical spaces but are well tolerated by users. Similarly, the structural aspects of virtual spaces that prove important for learning and navigation may differ from those of either a natural space or a carpentered city environment.

A third salient difference relates to the level of immersion of an observer at rest in the environment. Virtual environments are of much lower resolution than the human eye is capable of perceiving, and fields of view tend to be much narrower. The latter difference may prove to be particularly important, as studies of human spatial cognition report that spatial orientation (so called ambient space constancy) is particularly sensitive to peripheral stimulation (Goodale, 1996). So the relatively narrow field of view in all but the top-end CAVE displays may limit natural mechanisms for learning spatial environments, again necessitating alternative methods for wayfinding in those situations.

A final difference relates to the nature of the sensory feedback that results from movement in the virtual world. Observer movement in the physical world creates immediate full-field visual changes in the absolute and relative locations of objects (visual shear and motion parallax, respectively). Head and eye movements also create immediate feedback of a predictable nature. Unexpected changes in aspects of this interaction such as retinal blur reduction due to aliasing of VDT rasters (Macnik, Fisher and Bridgeman 1991) and greater than expected retinal shear due to misalignment of the optical and rotation axes of the eye (Hadani, Ishai, and Gur, 1980) may impact spatial orientation in unusual ways.

It is important to note that these effects are not independent. Differences between performance characteristics of display methods, regularity and complexity of environmental structure, and individual differences in perception, learning, and cognition may interact, creating a variety of performance breakdowns.

The impact of these environmental and the perceptual differences between virtual and real environments on higher-order tasks such as wayfinding is unknown. This study begins an attempt to characterize the ways in which performance on familiar tasks, such as finding a particular location in a space, will evolve over time in a relatively unstructured virtual environment. We can then examine whether it is possible to transfer spatial orientation methods from the natural to the virtual worlds, and perhaps understand the impact of complexity and regularity of structure. The primary goals of the present study were to test the effectiveness of our experimental design and Repeating and Wayfinding tasks,and to examine the value of active steering, compared to passive following in enabling VR users to develop a mental map of a VE. We also implemented and evaluated a form of "training wheels" that prevented trainees from straying too far from the desired path when actively following with the idea that this could provide the benefits of both active steering and passive following.

## Method

### Participants

Sixty-three paid volunteer participants were recruited through postings to local network newsgroups, posters placed around campus, and ads placed in campus newspapers. All participants were university students, with experience in using a mouse to interact with a graphic user interface. No limitations were placed on the age of the participants, and the participants spanned from the late-teens to mid-40's. Sex was not a factor in recruitment, though sex was balanced across all conditions. The field of study of the students was also not a factor, although computer science accounted for the majority of the participants (33).

### Equipment

The virtual world in the study was rendered on an Indigo2 Impact 10000 Silicon Graphics workstation. The software consisted of a custom system, FledermausVR, which presents a perspective view of the terrain, and a mouse-based flying interface (Ware and Fleet, 1997). The participants controlled their navigation with this mouse-based interface, which consisted of widgets that provide 6 degrees of freedom in movement. This existing system was written in C++ and OpenGL, and some alterations specific to the tasks of this study were added. The terrain was displayed non-immersively, on a standard 19-inch color monitor, with a resolution of 1280 pixels horizontally and 1024 vertically. Neither stereo nor head coupling was used. The point of view within the virtual world was updated dynamically, in real time, in response to the mouse actions. The participants were free to arrange the monitor and mouse to suit their preferences. The ambidextrous flying interface allowed both for right-handed and left-handed participants.

Procedure

The task that participants were required to perform was to navigate in a 3-D virtual environment by driving through the terrain from the start point of each path to the designated end-point target. These targets were colored cubes, with two of the four side walls removed to provide entry and exit points. The study proceeded through five distinct phases: *Familiarization, Testing, Learning, Repeating*, and *Wayfinding*.

In the *Familiarization* phase, participants were placed above a flat terrain, which contained two teapots sitting on the virtual ground (Figure 1). The purpose of this phase was to give participants an opportunity to become familiar with the experiment system, and to gain some practice in driving using the 6DOF mouse interface. To encourage this process, participants were asked to identify the flavor of tea in each teapot. This task required that they descend to ground level, approach the teapots, and circle them until they could read the labels on one side of each teapot.

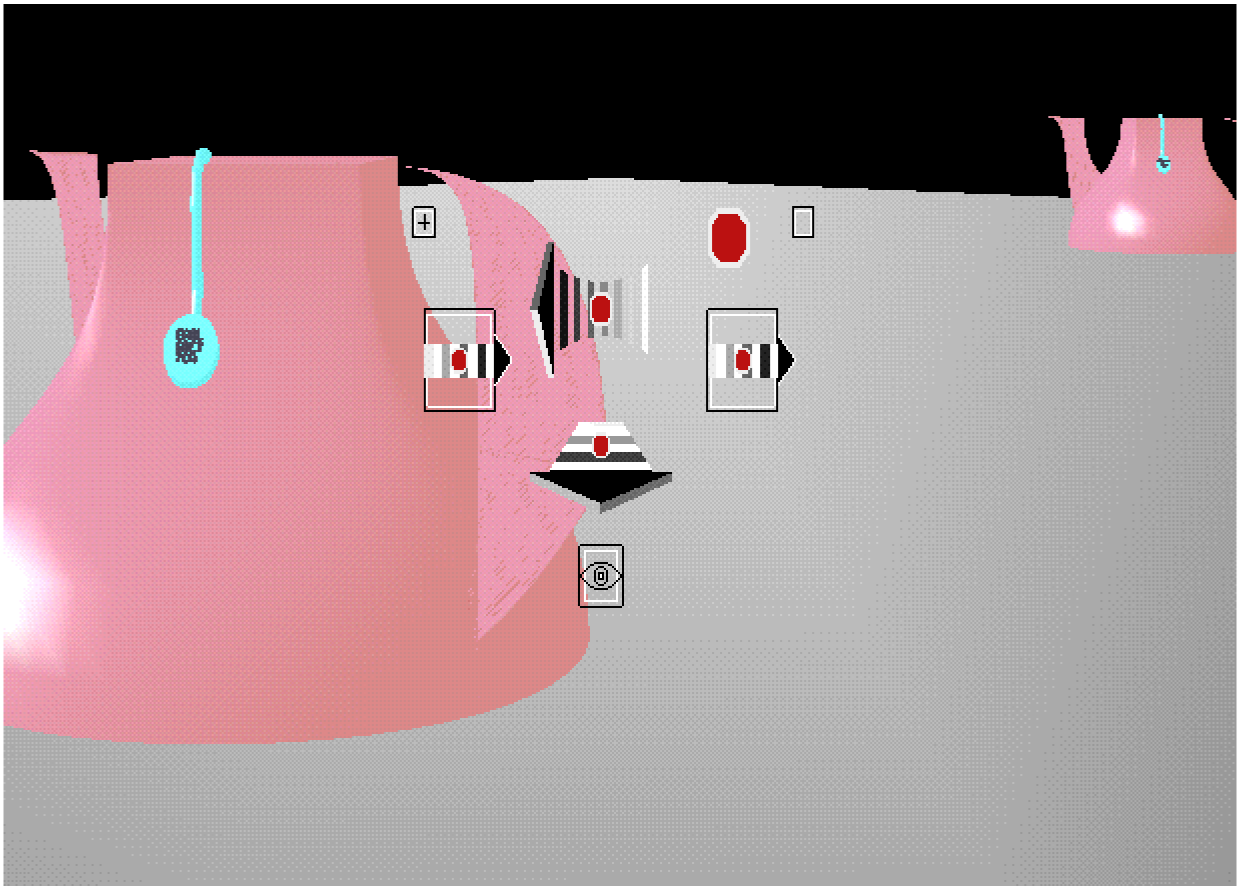

Figure 1: Screen shot of "Familiarization" phase (28K)

Upon completion of the Familiarization phase, participants proceeded to the Testing phase. The purpose of this second phase was to ensure a minimum level of comfort and competency with the interface before participants were allowed to advance to the main body of the experiment. This phase, therefore, served as a pre-test Driver's Exam; the participation of those participants who were not able to complete it within the time limit was terminated.

The Testing phase task (Figure 2) was to traverse twelve targets, scattered throughout another flat terrain. Two visual cues were provided, to ease this task. First, the targets were placed progressively farther along a blue trail, visible at all times, which ran along the ground. Second, the targets were colored such that traversing a target with its red wall on the right side indicated a proper direction along the path. Only one of the twelve targets was visible at a time, and constituted the "next target"; upon successful traversal, this target would disappear and the subsequent one, farther along the blue trail, would appear. The task tested the ability of participants to change elevation in the terrain, by varying the height at which the target was raised from the ground by its base. The trail also required that participants steer through both sharp and gradual turns to both the right and the left.

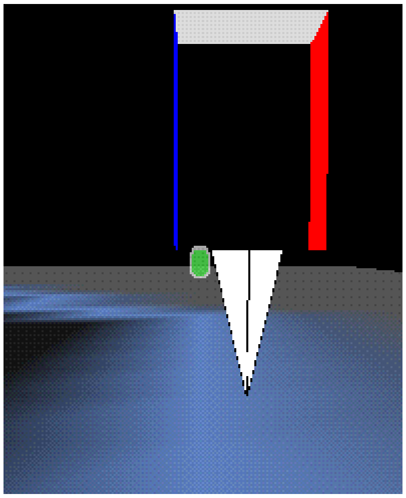

Figure 2: Screen shot of start of Testing phase. The first target is directly ahead.

The remainder of the experiment was conducted in a three-dimensional terrain, consisting of hills, plains, valleys, and mountain passes. The terrain was colored according to height, to better distinguish landscape elements from one another. Ten scattered locations within the terrain were identified with uniquely colored boxes, or stations. These stations were slightly larger than the targets of the Testing phase, and were set just above the ground. Half of the stations were exclusively Start stations (points of departure), from which trials were begun. The other half were End stations, destinations at which trials conclude. The task in each of the remaining phases was to drive from one Start station to a designated End station. Every trial required at least one and as many as three turns of at least 90 degrees. The ten stations were visible at all times, thereby becoming additional landmarks to the terrain, by which the participants might orient themselves. The height above the terrain to which participants could ascend was limited to approximately 75% of the height of the tallest mountain in the landscape.

The Following phase introduced the participants to this landscape and the stations. The task was to follow another vehicle (called the pilot vehicle) along a set of paths through the virtual world, travelling from a Start station to an End station in another part of the terrain. Upon completion, the next path was loaded, and the subject was automatically relocated to the next Start station. Familiarity with the terrain was developed through repeated visits to the same plains, mountain passes, and End stations. Participants were asked to observe the specific path traveled by the pilot vehicle, since they would subsequently be asked to duplicate it without the benefit of a pilot vehicle. They were also asked to observe the terrain as a whole, in order to be able to find new paths through it on their own.

Active and passive conditions. For this phase, each subject was randomly assigned to one of two main terrain-learning conditions, and to one of two sub-conditions within each main division. The main conditions were whether the subject was required to actively steer along the path of the pilot vehicle, or whether they were passively towed along at a fixed distance.

Within the Passive condition, the participants had no control over the path or speed of their vehicle. Participants were assigned to the Frontseat condition, in which the steering interface was visible, akin to watching the driver from the frontseat of a vehicle, or to the Backseat condition, in which the steering interface was absent, akin to sitting in the backseat of a vehicle, watching the terrain go by.

Within the Active condition, participants had to steer behind the pilot vehicle, but were not given complete freedom in the terrain. Their location and position relative to the pilot vehicle was monitored, and they were prevented from wandering too far from the correct path, falling too far behind, or passing the pilot vehicle. When participants exceeded these invisible boundaries, their heading was adjusted to match the correct path, and their velocity was reduced; persisting in leaving the correct path results in coming to a complete halt until corrective action is taken. Participants in the Active condition were randomly assigned either to the Tight Constraints condition (maximum deviation of 10% of the terrain before intervention) or the Loose Constraints condition (maximum deviation of 20% of the terrain before intervention). This effectively gave us two levels of active following, and it was hypothesized that the more active "Loose" condition might lead to better performance than the "Tight" condition in the final two phases of the experiment.

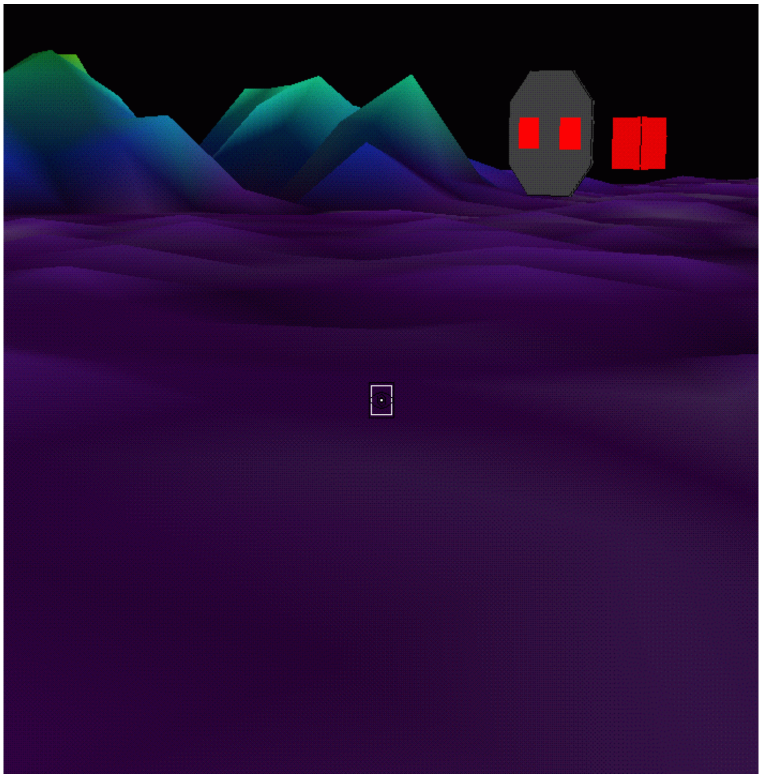

Figure 3: The Pilot Vehicle leads a subject through a right-hand turn.  One station is visible in the distance.

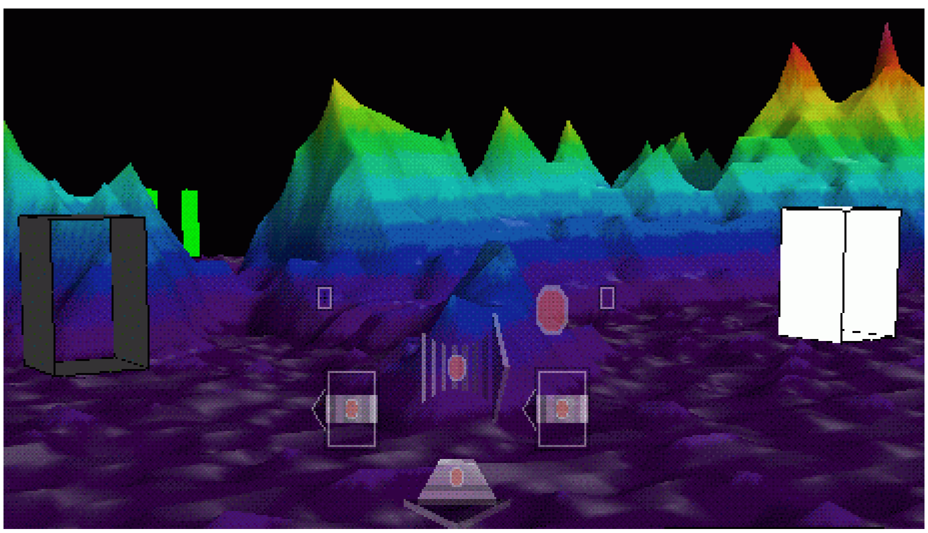

Figure 4: Screen shot of part of the terrain.  3 stations are in view.

In the Repeating phase, participants were asked to duplicate the paths traversed in the Following phase.  The pilot vehicle and all constraints were removed, and all participants (including those in the Passive conditions) had complete control of the vehicle. The computer recorded their times.

Finally, in the Wayfinding phase, participants were given a set of trials from known Start stations to known End stations, but for which they have not been shown an explicit path. They were then asked to find their own paths between the points of departure and destination. By this point, participants had seen each Start station and each End station at least four times during the Repeating phase. As in all phases, the height at which the vehicle could travel was constrained to prevent participants from gaining a "birds eye view" of the landscape.

To complete these five phases, participants were scheduled for a single session, two hours in duration. After a brief introduction to the experiment and the FledermausVR system, they advanced through the five phases described above. At the conclusion of the study, participants completed a brief questionnaire, assessing their subjective reactions to the tasks for the condition in which they were placed. Participants were also asked to locate the ten stations on a colored map of the terrain, as viewed from above.

## Results

Visual inspection of the results indicated that for each of the twenty paths (10 Repeating, 10 Wayfinding) there were a few outliers (response times that were much longer than for the other participants). These extreme scores were removed from the data set. In total, there were 32 scores deleted (17 from the Repeating paths; 15 from the Wayfinding paths); three of one participant's twenty scores were deleted, and two of two other participants' scores were deleted, the remaining 25 scores that were deleted were from 25 different participants; five scores were deleted from Repeating path 7, and four scores were deleted from Wayfinding path 2, no more than two scores were deleted from any other path.

Paths in this study differed in length. To compare performance across paths of different lengths, each participant's travel time for each path was divided by the path length (measured in points). This gave us an inverse measure of efficiency of travel, where low numbers indicate better performance. Five repeated measures ANOVAs were calculated on the resulting scores. The first two ANOVAs investigated differences between conditions on the Repeating paths. The second two ANOVAs investigated differences between conditions on the Wayfinding paths. The final ANOVA investigated differences between condition on all twenty paths (10 Repeating, 10 Wayfinding).

Repeating Paths

Two parallel repeated measures ANOVAs ($\alpha = .05$) were calculated with the active/passive learning condition as the between-subjects factor, and path (ten levels) as the within-subject factor. The dependent variable was the response time divided by distance score. In the first analysis, there were four conditions: Passive-Backseat, Passive-Frontseat, Active-Tight, Active-Loose. In the second analysis data were collapsed across the non-significant passive/active and frontseat/backseat conditions, leaving two conditions: Passive, Active.

In the first analysis, the main effect for path was significant, $F(9, 288) = 15.09$, $p < .001$, and the main effect for condition (four levels) approached significance, $F(3, 32) = 2.76$, $p = .06$. The interaction was not significant.

In the second analysis, the main effect for path was again significant, $F(9, 306) = 15.82$, $p < .001$, and the condition (two levels) x path interaction was significant, $F(9, 306) = 2.67$, $p = .005$.

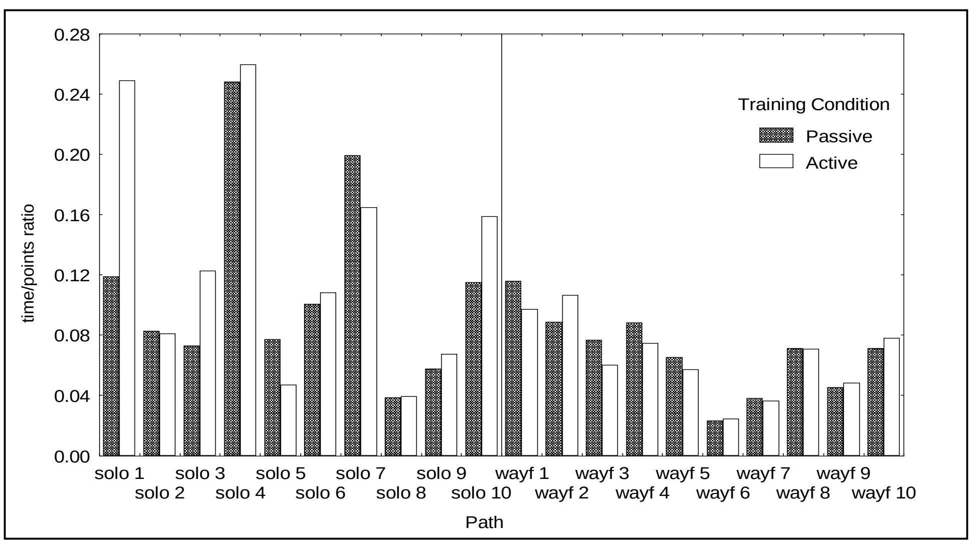


Figure 5: Performance of Active versus Passive learning subjects on path following for each path

In the near-significant main effect for condition (four levels) in the first analysis, the mean score for the Passive Backseat (.10) condition was significantly lower ($p < .03$) than the Passive

Frontseat (.14) and Active Tight (.14) conditions. The Active Loose condition (.13) did was not significantly different from any of the other three conditions.

In the second analysis, the main effect for condition was not present, but the condition (two levels) x path interaction was significant (Table 1). Tukey's Least Significant Differences (LSD) comparisons indicated that this interaction was the result primarily of differences in scores between paths in each condition. However, there were significant differences between the Passive and Active conditions on two of the paths ($p < .05$).

The main effect for path, present in both analyses, simply reflects the varying levels of difficulty in navigating from the start to end point on the ten paths.

<u>Wayfinding Paths</u>

Two parallel repeated measures ANOVAs ($\alpha = .05$) were calculated with passive/active learning condition as the between-subjects factor, and path (ten levels) as the within-subject factor. The dependent variable was the response time divided by distance score. In the first analysis, there were four conditions: Passive-Backseat, Passive-Frontseat, Active-Tight, Active-Loose. In the second analysis there were two conditions: Passive, Active.

In the first analysis, the main effect for condition (four levels) was significant, $F(3, 31) = 3.18$, $p < .04$, as was the main effect for path, $F(9, 279) = 15.82$, $p < .001$. The interaction was not significant.

In the second analysis, the main effect for path was again significant, $F(9, 306) = 15.82$, $p < .001$. There were no other significant effects with $\alpha = .05$.

In the significant main effect for condition (four levels) in the first analysis, the was mean score in the Passive Backseat condition (.09) was significantly ($p < .03$) higher (i.e., slower) than either the Passive Frontseat condition (.06) or the Active Tight condition (.06). The mean score for the Active Loose condition (.08) did not differ significantly from any of the other three conditions.

The main effect for path, present in both analyses, again reflects the varying levels of difficulty in navigating from the start to end point on the ten paths.

Comparing the Repeating and Wayfinding Paths

In the interest of comparing performance on the Repeating and the Wayfinding trials, two repeated measures ANOVAs ($\alpha = .05$) were calculated with condition as the between-subjects factor, and path (twenty levels) as the within-subject factor. The dependent variable was the response time divided by distance score. In the first analysis, there were four conditions: Passive-Backseat, Passive-Frontseat, Active-Tight, Active-Loose. In the second analysis there were two conditions: Passive, Active.

In the first analysis, the main effect for condition (four levels) was marginally significant, $F(3, 23) = 3.07$, $p = .05$, and the main effect for path was highly significant, $F(19, 437) = 18.64$, $p < .001$. The condition (four levels) x path interaction was also significant, $F(57, 437) = 1.45$, $p = .02$.

In the second analysis, the main effect for path was significant, $F(19, 475) = 17.61$, $p < .001$, as was the condition (two levels) interaction, $F(19, 475) = 1.68$, $p = .04$.

In the main effect for condition (four levels) in the first analysis, the only significant difference was between the Passive Backseat (.12) and the Passive Frontseat (.08) conditions. None of the other comparisons between conditions was significant with $\alpha$ set at .05.

In both analyses, the condition x path interaction was significant, but it was not possible to calculate any *post hoc* analyses due to the near singularity of the matrix. However, there was evidence of a practice effect (Figure 1), in which the mean scores for the wayfinding paths were lower (i.e., faster) than the mean scores for the repeating paths.

Discussion

In the Repeating task, subjects were asked to retrace a path they had followed earlier. Performance on 8 out of 10 paths did not show any effect of the experimental conditions. This could be a result of insensitivity of the dependent measure (i.e. the determining factors in task performance on those paths were not related to wayfinding). In other words, many of the paths may have been relatively easy to remember, and users' ability to wayfind was not tested. A second possibility was a lack of an effect of our experimental manipulations on users' ability to wayfind in this environment, which is considered less likely given the consistent effect of active/passive training condition for two of the paths.

For the paths that did show an effect, users were significantly slower in retracing the path when their first exposure to it was as a driver than when it was as a passenger. This is an unexpected result in light of everyday experience as a passenger versus driver in an automobile. Anecdotal reports typically agree that driving facilitates learning a route in the physical world better than being a passenger in the car. This is often attributed to increased attention to the environment that is required for effective driving. The more passive role of passenger does not demand that level of attention, and so does not trigger the formation of a mental representation of the space and the corresponding wayfinding ability. Our findings contradicted this common-sense view, and may be due to a variety of factors:

**Environmental structure explanation:** The unfamiliarity of the environment and lack of structure required users to expend a great deal more perceptual attention to understanding and remembering the surroundings than is the case in the familiar, structured world of city streets. This would suggest that a similar effect might be found in unstructured environments in the natural world.

**Attentional limitation explanation:** The lack of familiarity of the control interface may have facilitated this effect. The novel task of piloting a vehicle in a virtual world using a new interface is likely to require a great deal more attention than the overlearned task of driving an automobile. This would have required subjects to divide their attention between tasks, impairing the creation of a mental representation during the learning phase in the active driving (but not the passive passenger) condition. This possibility is supported by our finding that "Repeating" performance on the Backseat condition (no visible control interface) was significantly better than Frontseat (visible controls). The display may have drawn attention away from the landscape, inhibiting the formation of route memory.

**Perceptual limitation explanation:** The restricted views of the environment allowed by our single screen display did not extent to the retinal periphery. Peripheral retinal information has been found to be important in human visual orientation in a large number if studies, and the lack of static and optical flow (structure-from-motion) structure could have impacted learning. Follow up studies are planned using wide-field displays to test this hypothesis.

The Wayfinding trials did not show any effect of passive versus active learning, and the high level of overall performance in this phase suggests that participants found the task too easy.

A more difficult route might be more sensitive to differences in wayfinding ability between passive and active learning conditions. Similarly, no effect was found for Tight versus Loose coupling of drivers to pilot vehicle during the learning phase in either Repeating or Wayfinding tests. This may be due to the choice of levels of coupling, individual differences, or to a general insensitivity of the dependent measures.

In conclusion, the experiment described here examined the impact of active versus passive learning on route memory and overall mental map formation. The results were not in agreement with our predictions, and contradicted the common belief that active involvement in moving through a terrain promotes spatial memory (i.e. drivers have better memory than passengers). This points out the perils of extending beliefs based upon everyday experience in the physical world to virtual worlds, where perceptual, structural and task differences may lead to unexpected patterns of performance.

## References

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

Kosslyn, S. M. (1981). The medium and the message in mental imagery: A theory. Psychological-Review 88(1): 46-66.

Pylyshyn, Z.W. (1981). The imagery debate: Analogue media versus tacit knowledge. Psychological-Review 88(1): 16-45.

Hutchins, E. (1996). Cognition in the Wild. MIT Press, Cambridge MA.

Lynch, K., Banerjee. T., Southworth, M., (1995). City Sense and City Design : Writings and Projects of Kevin Lynch MIT Press, Cambridge MA.

Ware, C., and Fleet, D (1997) Context Sensitive Flying Interface.Proceedings 1997 Symposium on Interactive 3D graphics, Providence, RI, April, 127-130.

Hadani,-I.; Ishai,-G.; Gur,-M. (1980). Visual stability and space perception in monocular vision: Mathematical model. Journal-of-the-Optical-Society-of-America.70(1): 60-65.

Vinson, N. G. (1999). Design Guidelines For Landmarks to Support Navigation in Virtual Environments.CHI'99 Conference Proceedings.

Fisher, B.D., Agelidis, M., Dill, J., Tan, P., Collaud, G., & Jones, C. (1997). CZWeb: Fish-eye views for visualizing the World Wide Web. In M.J. Smith, G. Salvendy & R. J. Koubek *Design of Computing Systems: Social and Ergonomic Considerations: Volume 2, Advances in Human Factors/Ergonomics, 21B*, pp 719-722, Amsterdam: Elsevier.

Macknik, S. L., Fisher, B. D., & Bridgeman, B. (1991). Flicker distorts visual space constancy. *Vision Research. 31:12*, 2057-2064

Goodale,M.A. .(1996) One visual experience, many visual systems. In Attention and performance 16: Information integration in perception and communication. Attention and performance. (pp. 369-393). MIT Press Cambridge, MA