

Evaluating Motion Constraints for 3D Wayfinding in Immersive and Desktop Virtual Environments

Niklas Elmqvist
INRIA
Saclay, France
elm@lri.fr

M. Eduard Tudoreanu
University of Arkansas
Little Rock, USA
metudoreanu@ualr.edu

Philippas Tsigas
Chalmers Univ. of Tech.
Göteborg, Sweden
tsigas@chalmers.se

ABSTRACT

Motion constraints providing guidance for 3D navigation have recently been suggested as a way of offloading some of the cognitive effort of traversing complex 3D environments on a computer. We present findings from an evaluation of the benefits of this practice where users achieved significantly better results in memory recall and performance when given access to such a guidance method. The study was conducted on both standard desktop computers with mouse and keyboard, as well as on an immersive CAVE system. Interestingly, our results also show that the improvements were more dramatic for desktop users than for CAVE users, even outperforming the latter. Furthermore, the study indicates that allowing the users to retain local control over the navigation on the desktop platform helps them in familiarizing themselves with the 3D world.

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*Interaction styles*; I.3.6 Computer Graphics: Methodology and Techniques—*Interaction techniques*; I.3.7 Computer Graphics: Three-Dimensional Graphics and Realism

Author Keywords

navigation guidance, navigation aids, navigation assistance, motion constraints, evaluation

INTRODUCTION

Wayfinding in a 3D environment—planning and forming strategies where navigation is not the goal of the interaction but the means to solve some specific task [10]—requires a high level of spatial understanding of the structure of the environment [25]. However, this understanding is difficult to attain in an environment that the user is navigating on a computer [7, 9]. Often, the interaction mechanics of 3D navigation impose a high cognitive load on the users, leaving them few resources to learn the environment they are traversing.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2008, April 5–10, 2008, Florence, Italy.

Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00.

Immersive Virtual Reality systems with full 3D input and output devices, such as CAVEs and HMDs, can often remedy this problem, allowing users to better focus on familiarization with the 3D world instead of on the actual interaction [3, 18]. However, despite having been available on the market for considerable time, such specialized devices are unlikely to gain widespread use on consumer-level computers in the near future. For example, mouse and keyboard are the standard input devices even for highly three-dimensional computer games. Therefore, we must find ways to improve navigation even for basic desktop computers.

One such approach that has been introduced in the literature is navigation guidance using motion constraints. This method offloads some of the effort of 3D navigation from the users by partially controlling their movement, freeing them to learn the environment instead. In other words, the idea is to actually help the user by *taking away* some of the bewildering freedom that full 3D navigation offers. In this paper, we combine these existing methods into a unified strategy and investigate its benefits by comparing it to manual navigation. Our hypothesis is that navigation guidance using motion constraints will improve performance and memory recall, particularly for desktop computers with mouse and keyboard as input devices.

The guidance method should have the following requirements for it to be effective:

- *global coverage*: all of the important landmarks in the environment should be presented to users, preventing the users from becoming lost or stuck in a subset of the world;
- *continuous motion*: the movement should be smooth and continuous to promote environment understanding; and
- *local control*: local deviations must be allowed to make the user an active participant (as opposed to a passive recipient) as well as to promote landmark recognition.

We present a guidance method fulfilling these requirements that makes use of a pre-defined tour of the world, similar to [13, 28]. The tour can either be built manually by a human designer, or automatically generated by some algorithm (such as [1]). The tour augments the standard navigation controls, essentially “holding the hands” of the users as they traverse the world. Depending on the level of interaction desired, we can impose constraints on the path, speed, deviation, and camera direction.

We have conducted a controlled experiment involving navigation and memory recall tasks on both immersive CAVE and standard desktop platforms. The purpose of this study was to investigate the benefits of navigation guidance for both of these platforms.

The application domains of navigation guidance are many and varied: it can be used for visual storytelling when introducing a new 3D environment, familiarizing a 3D modeler or designer with an unknown or forgotten project, or presenting all the relevant information in a visualization space, etc. According to the requirements above, the method can be configured to either be unobtrusive, merely nudging the user in the right direction, or take full control of the user's movement through the world.

This paper is structured as follows: We begin with a review of existing work on 3D navigation in general and navigation assistance in particular. We present our unified 3D navigation guidance method. We then describe our user study and the results, and finish the paper with discussions and conclusions of our findings.

RELATED WORK

Effective navigation and wayfinding in 3D environments is a well-known and well-researched problem. In the following sections, we give a representative sampling of some of the existing work in the area.

Basic 3D Navigation

The flying, eyeball-in-hand, and scene-in-hand metaphors [26, 27] constitute perhaps the most basic 3D navigation techniques, and form the baseline for the more sophisticated techniques following them. Bowman et al. [5] present a taxonomy of first-person motion control techniques for manual viewpoint travel that is useful for evaluating these methods.

Mackinlay et al. [17] describe a method of logarithmically controlling the viewpoint speed while moving through a 3D world to allow for rapid motion over large distances, yet slows down when approaching the target. In related work, Song and Norman [23] propose a set of non-linear motion control techniques for intuitively traversing virtual environments. The work of Tan et al. [24] on a moded navigation technique is interesting, not only for the fact that it contextually combines two different motion control techniques (flying and orbiting), but also that it couples the speed of movement to the height and tilt of the camera to smoothly support both local detail views and global overviews.

Guided Navigation

Visual aids can be used to great effect for guiding 3D navigation. Static or interactive maps are the obvious examples, and can often be combined with waypoints and direction arrows to indicate where the user is expected to travel.

Chittaro and Burigat [7] present an array of different compass-like navigation widgets for helping the user to find important objects and places in a virtual environment. Trails [21] help users utilize previous exploration to improve their current search. Path drawing [16] lets the user draw an intended path directly on the 2D view of the world to aid navigation.

Virtual guides are physical avatars in the 3D world that help direct users where they should go. One notable such technique is the virtual guide of Chittaro et al. [8] that the user must follow actively. The guide's path through the world is automatically computed using an algorithm operating on a 2D occupancy matrix. Yoon and Maher [30] present another virtual guide approach based on a swarm algorithm that dynamically explores a virtual environment.

Motion Constraints

Another powerful class of navigational aids is motion constraints, i.e. different ways of constraining the user's movement while traveling through a virtual environment. Techniques in this class can have varying degrees of obtrusiveness, from merely nudging the user in the right direction to constraining or controlling the viewpoint completely.

Wernert and Hanson [28] present a taxonomy of assisted navigation through motion constraints, and also discuss a "dog-on-a-leash" approach to guidance through a 3D world. This approach is similar to the "river analogy" introduced by Galyean [13], where the viewpoint is tethered to a vehicle following a path through the virtual environment and some degree of control is retained by the user. While these techniques fulfill our requirements from the introduction, we are looking for an even more general formulation of navigation guidance using motion constraints. Also, neither of the two include an empirical evaluation of their effectiveness for real user tasks.

Finally, some techniques assume full control of viewpoint motion, sometimes even moving the gaze of the user in the desired direction. By reducing the freedom of the user, navigation and wayfinding can be simplified, and eliminate the need for expensive features such as collision detection. Examples of this approach include that of Hanson and Wernert [14], who employ invisible surfaces to constrain user movement, and of Hong et al. [15], who provide an interactive navigation where care is taken to avoid collisions with the environment.

The Way-finder system presented by Andújar et al. [1] algorithmically computes an exploration path through a 3D environment in an off-line phase, ensuring that the whole environment is visited. The user is then taken along this path in an interactive phase.

CubicalPath [4] uses an off-line phase to build a volume representation of the 3D world and then employs dynamic potential fields for guiding exploration at run-time. In contrast to the mostly path-based techniques described above, this results in semi-constrained motion, where the user is nudged in the right direction to visit all targets in the world.

Motion constraints may also have detrimental impact on 3D navigation performance: Witmer and Singer [29] discuss the value of user control for navigation efficiency and presence.

DESIGN FRAMEWORK: WAYFINDING IN 3D

Wayfinding is the cognitive task of planning and forming navigation strategies for an environment prior to executing them; e.g. the actual navigation is not the goal of the inter-

action but the means to solve some specific task [10]. Because the environment is typically larger than can be seen from a single viewpoint, the wayfinding task is conducted on a mental representation of the environment, often called a *cognitive map* [6, 25]. If the user lacks an accurate cognitive map of the environment, wayfinding performance will suffer.

The objective of our work is to support wayfinding by helping the user build and maintain a correct cognitive map of the 3D world. The approach in this paper is based on actively guiding the user through the world, akin to taking a visitor on an initial sightseeing tour of your city or neighborhood to give them a basis for navigating on their own.

However, wayfinding in a virtual 3D environment has additional challenges beyond those for the physical world:

- The user's view of the virtual world lacks the visual fidelity of the physical world and is often not immersive;
- The controls for exploring the virtual world, typically standard mouse and keyboard on consumer-level computers, are artificial and poorly mapped to the task; and
- The virtual world lacks many of the additional sensorial stimuli of the physical world [7], including as subtle cues such as touch, temperature, wind, balance, etc.

All three of these factors combine to cause a high cognitive load on users trying to orient themselves in a complex 3D virtual environment. This means that users must expend considerable effort just to master the mechanics of the navigation, leaving few resources for building and updating their cognitive map [2]. Put differently, the user's *situational awareness* is decreased. While some of this cognitive load can be reduced somewhat through the use of immersive 3D input and output devices (such as in Virtual Reality environments), this kind of specialized hardware is currently expensive and uncommon on the consumer-level market.

Cognitive Maps

For this purposes of this paper, we will use a very simplistic model of cognitive maps, see Tolman [25] or Chase [6] for more sophisticated models. Basically, we define a cognitive map as a set of *landmarks* and their *spatial relations*. Consider a tourist visiting Paris for the first time. Figure 1 shows an example of how the tourist's cognitive map might look after a busy day of sightseeing—a collection of landmarks and some knowledge of how to go between them.

In our context, cognitive maps can be used in two primary ways:

1. *orientation*: locating a place in the real environment in the cognitive map (and vice versa); and
2. *wayfinding*: planning a navigation strategy for a specific task (such as visiting a landmark).

In our example, orientation corresponds to being dumped in the middle of Paris (as in, emerging from a Métro station) and having to figure out where, say, your hotel is located in relation to your current position. Analogously, wayfinding is the task of planning how to get to there from your position

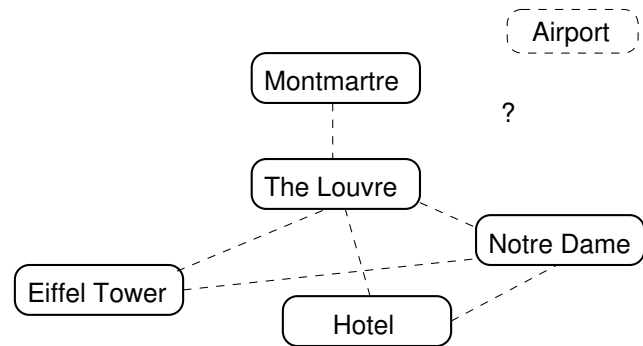


Figure 1. A tourist's cognitive map for downtown Paris.

once you have oriented yourself.

Naturally, cognitive maps may be more or less accurate and complete. Actual landmarks or spatial relations between them may be missing or incorrect. Furthermore, the user's knowledge of the appearance of certain landmarks may be incomplete or inaccurate, causing identification errors. These may all have repercussions for the two high-level tasks listed above.

Going back to our example, the tourist's cognitive map may be incomplete and potentially inaccurate; a nighttime taxi trip from the airport means that the tourist has no clear idea where the airport really is. However, as the tourist spends more time in Paris, this cognitive map will gradually become more and more refined and closer to reality.

Supporting Cognitive Map Building

It is clear that simply reducing the cognitive load of 3D navigation will support users in building and updating their cognitive maps by allowing them to spend less effort on the mechanics of the navigation, as outlined above. However, there are additional ways to guide 3D navigation based on our model of cognitive maps above. Below we derive the requirements listed in the introduction in greater detail.

Global Coverage

Obviously, for a user's cognitive map to be complete, the user must be exposed to the whole environment, or at least all of its important landmarks. In particular, users often become lost or stuck in one particular part of a large environment. This may happen when users are either unable to visit the whole environment due to disorientation or physical barriers (such as geometry collision detection), or mistakenly believe that they have seen the whole environment already.

Continuous Motion

From the model of cognitive maps in the previous section, it is clear that spatial relations are vital for users to be able to efficiently correlate landmarks in 3D space. Therefore, a navigation guidance method needs to be based on smooth and continuous motion that allows users to correctly build and update the spatial relations in their cognitive maps [5]. One important detail here is that the guidance method should avoid passing through any impenetrable 3D geometry—a

process known as “ghosting”—since this may be very confusing and jarring for the user.

Local Control

Finally, navigation guidance should never take full control over the movement in 3D space from the user. There are two reasons for this: First, if the guidance method is controlling the navigation without the user having any influence, the user is reduced to a passive recipient instead of an active participant, causing slower and less accurate understanding [19]. Accordingly, the user’s efficiency at cognitive map building will decrease.

Second, by allowing for some local control over the 3D motion, the user will be able to better control how much time to devote for studying each landmark in the 3D environment, possibly from different angles, thereby increasing recognition of the landmarks.

NAVIGATION GUIDANCE USING MOTION CONSTRAINTS

We present a method for 3D navigation guidance with dual objectives: to help first-time users immediately discover all of the landmarks in the world, as well as to support them in building an accurate cognitive map of the world for future use. The basic idea is simple: 3D navigation is complex, so we employ motion constraints to reduce some of this complexity and to achieve both of our objectives.

Our approach is based on constraining the user’s motion to a predefined path that tours the environment. In addition to taking away some of the bewildering freedom that causes the high cognitive load in the first place, this approach also has the benefit of supporting the user maintaining an accurate cognitive map if the path is designed to visit all of the important landmarks of the environment.

Tour-based Motion: Ensuring Global Coverage

Our technique constrains navigation to a pre-defined grand tour of the world. The tour is either created manually by a human designer or generated automatically by an algorithm. As argued by Wernert and Hanson [28], this kind of guidance ensures that the users will see all of the important landmarks the designer has created in the world. Furthermore, tour-based motion also eliminates the need for collision detection, which can be computationally expensive.

Spring-based Control: Allowing for Local Deviations

Our guided exploration technique uses a spring-like umbilical cord that connects the viewpoint to the grand tour. This allows the user to locally deviate from the pre-defined path as far as the virtual spring will allow it. After the user is done, the spring will smoothly return the user to the path.

See Figure 2 for an overview. Tables 1 and 2 summarizes the navigation controls in our implementation. Depending on the level of interaction desired, we can impose constraints on the following properties:

- **Speed.** Movement along the tour can either be computer-controlled or user-controlled.

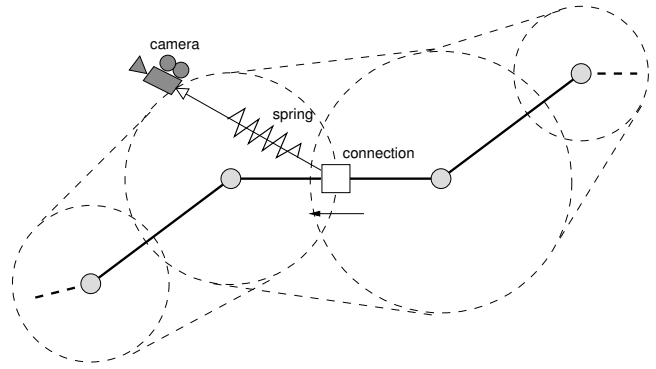


Figure 2. 3D navigation guidance (circles show free space).

- **Viewpoint direction.** The direction of the camera can either be slaved to the direction of movement, fixed to follow the currently closest landmark, or fully user-controlled.
- **Local deviation.** Using a simple interaction technique, the user can smoothly zoom the viewpoint forward or backwards in the direction of movement to the full extent of the virtual spring.

Smooth Animation: Supporting Continuous Motion

Since we use tours and not a set of simple viewpoints in space to ensure global coverage, we can also fulfill our second requirement by smoothly animating the user’s movement along the tour. The grand tours accepted as input are generally discrete keyframes, and so we fit Hermite curves [12] to these points to further smooth the movement through the 3D space.

Depending on whether the application supports collision detection or not, the viewpoint may either collide when it comes into conflict with world geometry, or it may float through the geometry as if it was not there. These two events, called collisions and ghosting, are often disorienting to the user, and is typically a major complaint when exploring a 3D world. We avoid these occurrences by computing the amount of free space around the tour in all points and constraining the full length of the umbilical cord to this value. This ensures a smooth and continuous ride through the environment with no jarring stops or confusing ghosting.

| PC Input | Interaction |
|-------------|---------------------------|
| LMB + mouse | pan view (pitch and yaw) |
| MMB | zoom forward |
| RMB | zoom backward |
| Up key | increase speed along tour |
| Down key | decrease speed along tour |

Table 1. Control interface for navigation guidance with mouse and keyboard controls. (L/M/R MB = left/middle/right mouse buttons.)

STUDY METHODS

We conducted a controlled experiment to study the efficiency of our technique for 3D navigation guidance using motion

| CAVE Input | Interaction |
|--------------------|-------------------------|
| Left/right ball | rotate view (yaw) |
| Wand + up ball | zoom forward |
| Wand + down ball | zoom backward |
| MWB + up/down ball | change speed along tour |

Table 2. Control interface for navigation guidance with wand control. Wand consists of a 6DOF tracker, a four-position ball/joystick, and three buttons. (L/M/R WB = left/middle/right wand buttons.)

constraints for cognitive map building. This section summarizes our methods.

Predictions

The basic premise of this research is that guiding the user in exploring a 3D world will promote cognitive map building and thus increase the user's future efficiency in solving wayfinding tasks compared to unguided navigation. We also predict that fully constraining the movement of the viewpoint will reduce the viewer to a passive recipient instead of an active participant. Accordingly, perception of the world as a whole will suffer even if the user is shown the important landmarks by the guidance technique. More specifically:

P1: Navigation guidance will result in better wayfinding performance than free navigation.

P2: The presence of user control will improve subject familiarization of the environment.

Furthermore, the input and output limitations of desktop-based 3D navigation means that much of the user's attention is devoted to the mechanics of moving around. Therefore, we also add a third prediction:

P3: Navigation guidance will benefit desktop platforms more than immersive platforms.

Participants

We recruited 35 subjects for this study (6 female, 29 male). 16 (4 female, 12 male) participants participated in the desktop part and 19 (2 female, 17 male) in the immersive one. The two groups were separated geographically because the two platforms were set up at different universities. The subjects were all undergraduate and graduate students from the engineering programs at our universities. Ages ranged from 20 to 40 years. All participants were screened to have at least basic computer skills, were not color blind, and had normal or corrected-to-normal vision.

Apparatus

The desktop-based version of the experiment was conducted on an Intel Centrino Duo laptop computer equipped with 2048 MB of memory running the Microsoft Windows XP operating system. The display was a 17-inch widescreen LCD display running at 1920×1200 resolution and powered by an NVIDIA Geforce 7800 GO graphics card. Input devices were a standard Microsoft 3-button optical mouse and the laptop's keyboard.

The immersive version of the experiment was conducted on a three-sided CAVE environment consisting of front, right

and floor display walls. Each wall was 10×8 feet powered by a Christie DLP projector with a resolution of 1280×1024 pixels. The CAVE was run by three dual-processor Intel Xeon 3.06 GHz personal computers with 3Dlabs Wildcat4 7210 graphics card, one computer for each display wall. Each wall provided an active stereoscopic image using CrystalEyes shutterglasses connected to a six degree-of-freedom Flock of Birds tracker. The input device was a three-button wand, also tracked by the Flock of Birds.

Experimental Conditions

We included three factors in our study: platform, navigation method, and scenario.

Platform

Since we hypothesized that navigation guidance would have a different impact on desktop-based as well as immersive platforms, we included this as a factor in our study: a *desktop* condition for a standard PC, and *immersive* for a CAVE system.

Navigation Method

The navigation method employed was one of the three in the list below:

- *Free.* First-person 3D navigation with no guidance.
- *Follow.* Passive tour following with full guidance except for camera orientation.
- *Spring.* Full navigation guidance with user-controlled movement, deviation, and camera orientation.

Scenario

The four different scenarios employed in the experiment were designed to depict typical usage situations of 3D worlds and 3D navigation using both abstract as well as realistic environments. Subjects were given a briefing of the scenario prior to starting each scenario run. The pre-computed tours were designed to pass targets within a distance of 5% of the world size.

Below follows a short description of each scenario, including target density in relation to the total number of objects:

- *Outdoor.* Large-scale outdoor world with a rescue mission scenario where the user was asked to identify helicopters, cars, and fire hydrants (Figure 3). [15% density]
- *Indoor.* Maze-like single-floor indoor environment representing a warehouse where the user was looking for office chairs, sofas, and floor lamps (Figure 4). [10% density]
- *Infoscape.* Abstract information landscape for a hypothetical 3D file browser where the subject was preparing for writing a report by looking for Word, PDF, and Excel files. (Figure 5) [5% density]
- *Conetree.* Abstract conetree [20] visualization for the organization hierarchy of a company where the subject looked for leaf nodes of specific colors. (Figure 6) [5% density]




Figure 3. Targets and environment for the outdoor scenario.

As can be seen from the description, the former two are real-world scenarios, whereas the latter two are abstract.

Study Design

The experiment was a $2 \times 3 \times 4$ mixed design on the following factors (WS for within-subjects, BS for between-subjects):

- Platform: *desktop, immersive* (BS)
- Navigation method: *free, follow, spring* (BS/WS)
- Scenario: *outdoor, indoor, infoscape, conetree* (WS)

The scenario was counterbalanced for all groups, and the navigation method was counterbalanced for the CAVE group (whereas it was between-subjects for the desktop group). The difference for the navigation method for the two groups was due to local variations at the two different universities the study was performed.

We collected the following metrics for each trial:

- Recall distance*: average distance between placed landmark and its actual position (phase II);
- Evaluation error*: average number of missed landmarks divided by the total number of landmarks (phase III); and
- Time per target*: average time for finding each landmark (phase III).

Procedure

Each test session lasted approximately one hour, and started with an introduction and a training session to allow the subject to get familiarized with the test application, the interface, and the test procedure. After the subjects indicated that they were satisfied, we proceeded with the actual trials.

The actual test consisted of four trials corresponding to the four scenarios. For the desktop group, each participant was randomly assigned to a navigation method. For the immer-




Figure 4. Targets and environment for the indoor scenario.

sive group, the participant was taken twice through each navigation method, and thus did two additional trials—the last two were repetitions of the first two scenarios, but with a different navigation method.

Each trial was tested in a three-phase sequence: (I) familiarization, (II) recall, and (III) evaluation. The three phases were designed to correspond to the formation of a cognitive map, the recollection of the map (short-time memory), and the application of map knowledge in the same environment. In the first phase, users were given the scenario world and were allowed to familiarize themselves with it for five minutes. During this phase, the guidance method selected for the user was active. The subject was given a reference card with pictures of three types of landmarks relevant to the actual scenario that he or she should be looking for. An overhead map of the world with the user's own position and location marked was available in the upper left corner of the display.

In the recall phase (II), the subject was shown a full-screen overhead map of the world and was asked to place as many instances of two of the three target landmarks as they could remember. There was no time limit in the recall phase.

Finally, in the evaluation phase (III), subjects returned to the 3D world and made use of the cognitive map developed in phase I to collect the third type of target landmark, the types not asked for in the recall phase (II). Here all subjects were forced to navigate freely with no guidance support. Collecting an object was done by approaching to within a distance of 5% of the world scale and pressing a key. This removed the object from the landscape. The miniature overhead map was available in this phase as well. When subjects decided that all targets had been found, they were able to end the trial (stopping the time).

The subjects did not know in advance which two of the three target types they would be asked to place in the second phase, nor which to collect in the third.




Figure 5. Targets and environment for the infoscape scenario.

RESULTS

The two most important factors in the analysis were Navigation method and Platform. Neither of the two factors, when considered separately, was found to be significant in determining the results. Nonetheless, there was significant interaction between the two factors for recall distance (a) and evaluation error (b) ($F_{2,34} = 12.09, p < 0.001$ and $F_{2,34} = 6.97, p = 0.042$, respectively). For time per target (c), the interaction was very weak to non-existent ($F_{2,34} = 2.57, p = 0.092$).

Navigation Method

For free navigation, the CAVE platform turned out to be both more accurate and quicker on average than the desktop platform; Figures 7 and 8 show the averages for each of the conditions. Only the recall distance (a) turned out to be significant ($F_{1,59} = 9.60, p = 0.003$) for this comparison.

On the other hand, when all subjects were given access to the “spring” navigation guidance technique, results for the desktop platform condition improved considerably to surpass the CAVE platform on all metrics. These were also all significant differences down to $p < .05$. When comparing the desktop condition with navigation guidance versus the CAVE condition with free navigation, the desktop participants again performed better both in terms of efficiency and accuracy. These differences were significant for recall distance ($F_{1,55} = 6.53, p = 0.013$) and time per found target ($F_{1,55} = 5.65, p = 0.021$), but not for errors.

Immersive Platform

For the immersive platform, recall distance (a) was strongly dependent on the navigation type ($F_{2,36} = 5.85, p = 0.0063$). A post-hoc Tukey test shows that there was a statistically significant difference between “follow” and “free” (adj $p = 0.039$), and between “free” and “spring” (adj $p = 0.0069$), yet no significant difference was found between the two tour-based navigation methods. In the CAVE, the evaluation error




Figure 6. Targets and environment for the conetree scenario.

was lowest for “free” navigation even though the time spent marking the targets was the shortest when compared to the other two methods “follow” and “spring” (see the blue bars in Figure 8).

None of the other dependent variables exhibited a significant variation with the navigation method. Specifically, there was no difference in time performance between methods.

Desktop Platform

For the desktop platform, the recall distance (a) was a significant factor ($F_{2,13} = 5.84, p = 0.016$). Post-hoc analysis reveals that only “free” and “spring” navigation are different (adj $p = 0.013$). Subjects were more accurate in placing landmarks in the recall phase for spring-based navigation guidance than for the free flight. This shows that spring-based guidance gave significantly more accuracy than both other methods for the desktop condition.

For the time per target (c), results confirmed that subjects with navigation guidance are more efficient than those without (red bars in Figure 8). This difference was also significant ($F_{2,13} = 7.12, p = 0.0081$). The post-hoc analysis shows that “follow” and “free” navigations were different (adj $p = 0.025$), as were “free” and “spring” (adj $p = 0.013$). Thus, there is a definite time advantage for any type of navigation guidance compared to free-flying.

DISCUSSION

We can summarize the results from the study in the following main findings:

- Unaided (free-flight) navigation generally resulted in better performance on immersive than desktop platforms;
- Guiding 3D navigation helped participants on the desktop platform outperform all CAVE participants; and
- Allowing local deviations for navigation guidance caused improved correctness for desktop platforms.

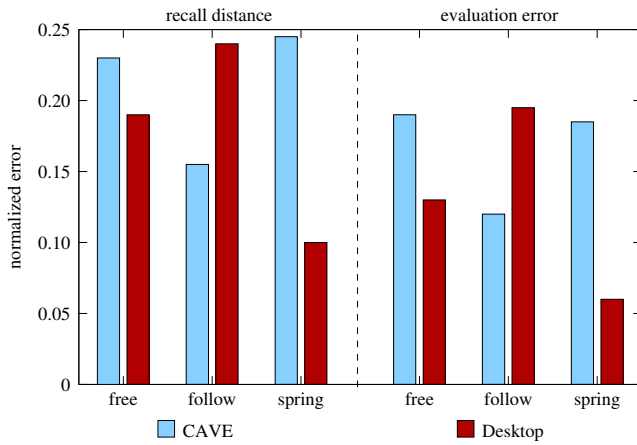


Figure 7. Recall distance (a) and evaluation error (b) normalized between 0 and 1. This figure shows that “free” navigation is the best for the CAVE, while “spring” is the best for the desktop.

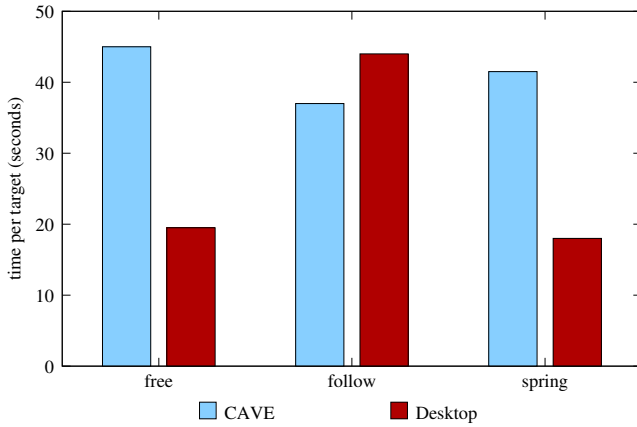


Figure 8. Average time to find a target landmark (c). Note that “free” navigation was the fastest on average for the CAVE, and that “spring” was the fastest for the desktop.

In the next sections, we will try to explain and generalize these results, discuss the limitations of the method, and describe our future work.

Explaining the Results

The fact that CAVE participants performed better than desktop participants with unaided navigation confirms one of the premises of this paper: that the cognitive effort of navigation in a 3D world on a desktop computer with a mouse and keyboard has a serious impact on wayfinding. For the CAVE platform, the immersive display and 3D input devices allow users to better focus on orienting and familiarizing themselves in the 3D world, enhancing their situational awareness and understanding of the 3D world.

Our results also show that introducing navigation guidance (both spring-based and follow-based) caused better performance for the desktop platform, significantly better than for the CAVE platform with and without navigation guidance. In other words, prediction P1 is partially confirmed for the desktop. We believe that this is because navigation guidance

reduces the cognitive effort of wayfinding with a mouse and keyboard. In other words, some of the advantages that an immersive platform brings for 3D navigation can be compensated for by introducing computer-supported navigation guidance that allows users to focus less on the mechanics of the interaction. This result is also in line with the findings of Schulze et al. [22]. Their experiment compared CAVE and fish-tank VR of a scientific visualization for a 3D marking task that is similar to phase III in our trials. Their results show that fish-tank VR (i.e. desktop-based) users actually perform better than immersive users in some conditions.

For the desktop platform, allowing for local deviations from the tour using the spring-based method did cause participants to perform better than if they were passively following the tour. This did not hold for the CAVE platform, however, where it only seemed to matter whether or not the user is completely free. This partially confirms prediction P2.

One possible explanation for the local control prediction not being proven in the general case might be that the subjects in the passive follow group were not truly passive since they were given a very specific task when familiarizing themselves with the 3D world. Therefore, they performed better than they might have done without this knowledge. However, our pilot testing showed that the alternative, i.e. not telling the subjects which kinds of landmarks to look for, was simply not feasible for the high-detail scenarios we used in the study.

It is also interesting to note that CAVE subjects performed *worse* using navigation guidance than free, unaided navigation. In a sense, this confirms prediction P3: navigation guidance certainly was of more benefit to the desktop platform. The reason for this may be that in the immersive case, the user is in a much better situation to perform unaided 3D navigation, and so being limited to a fixed tour may be perceived as annoying and detrimental to productivity. Another possible explanation is that it may feel unnatural to be dragged along in a CAVE given that an immersive environment has a heightened sense of presence.

Generalizing the Results

Generalizing these results comes down to many factors. We believe that the four scenarios we employed as part of this study constitute a fair cross-section of the kind of 3D environments users are typically asked to navigate in. For the future, it would be interesting to perform a similar study in a real computer game 3D environment.

The local differences between the desktop and immersive groups in the design of experiment may have had an impact on our results, although no proof was found in an analysis of the first trials. Specifically, the navigation method was between-subjects for the desktop group, and within-subjects for the CAVE group. This may have exposed the participants to different learning effects as the experiment went on. However, analyzing the first trials for the CAVE group showed similar results to the overall result, suggesting these results generalize. Moreover, both designs were counterbalanced, and the users received proper training for all navigation methods.

Our choice of participants may also have an impact on the study: all participants were students from our local universities. Accordingly, most participants were common computer users (more than 20 hours/week), and many had extensive 3D game experience. However, subjects were randomly assigned to navigation methods, so this fact should have no direct impact on the general finding, but definitely on the time and correctness measures we collected.

We designed our three-phase experiment based on our model of cognitive maps discussed in the design framework in the beginning of this paper. Given no previous knowledge of the environment they encountered, users were given five minutes to build an initial cognitive map that they could later use for orientation and wayfinding. However, our experimental design has no way of checking retention over time; for that, we would also have to have our participants come in a week or a month later. We did not have the resources for this in the current work, but it is an important component to be able to fully generalize the results.

That local control has a positive impact on the accuracy of the user's cognitive map seems to be at least intuitively generalizable. The analogy is driving a car yourself under someone's directions as opposed to being a mere passenger in a car driven by the knowledgeable person—remembering the way is much easier in the former case. More stringent tests are needed for studying this effect in future experiments.

Limitations to Navigation Guidance

While navigation guidance using motion constraints have many advantages, it is far from a universal solution that applies to all situations. As we have already seen, subjects actually performed *worse* using navigation guidance on CAVE platforms than using unaided 3D navigation. This suggests that we should be careful to employ navigation guidance only where it actually helps the users.

Regrettably, three of the participants in the user study (one CAVE user) became motion sick; one still finished the study (desktop), the other two canceled and their results are not reported at all in the paper. An interesting observation is that all of these participants were assigned to the passive tour following group. A possible explanation may be that users that have no control over their movement run a greater risk of becoming sick, somewhat akin to how people who are prone to motion sickness while riding cars typically only get it when they are passengers and not driving themselves.

The navigation guidance implementation that we tested in the study currently does not constrain the view direction, meaning that the user may very well miss a landmark entirely by looking the wrong way at the wrong time even if the landmark is visible from the tour. This can either be solved by controlling the view direction and tracking exactly which landmarks have been seen so far by the user, or by indicating which landmarks are currently visible but off-screen.

Finally, the effectiveness of navigation guidance using motion constraints is ultimately at the mercy of the quality of the underlying tour. Building a smooth and coherent tour that traverses the whole 3D environment in as short a dis-

tance as possible can be a challenging and time-consuming task for a human designer. For dynamically-generated 3D worlds, it is impossible. Automatic tour-generation approaches [1, 11, 14, 15] may come in useful for such situations.

Future Work

Our future work will focus on evaluating navigation guidance in real-world 3D environments such as computer game worlds. This will also allow us to study temporal aspects: regular players of online 3D multiplayer games have spent many hours in these 3D environments, and thus have highly detailed cognitive maps of them.

CONCLUSIONS

We have presented a method for navigation guidance in the exploration of general 3D environments intended to both promote the user's building of a cognitive map of the environment as well as to improve visual search task performance. The technique uses motion constraints to guide the user on a pre-defined tour through the environment while still allowing users to control their speed as well as to deviate locally from the tour.

To test these hypotheses, we conducted a user study involving both desktop-based as well as immersive VR. Interestingly, the results indicate a strong dependence of the effectiveness of a guidance method with the type of platform being employed. Overall, the best method was the new navigation guidance for a desktop platform, followed by free flying in an immersive platform. One interpretation of these results could be that navigation guidance allows users—particularly desktop-based ones—to focus less on the mechanics of 3D navigation using 2D input and output devices, helping them in building a more accurate cognitive map of the 3D world.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous participants who took part in our user studies. Thanks to Klaus Mueller of Stony Brook University for valuable discussions regarding the interpretation of our results.

REFERENCES

1. Carlos Andújar, Pere-Pau Vázquez, and Marta Fairén. Way-finder: guided tours through complex walkthrough models. In *Proceedings of Eurographics*, pages 499–508, 2004.
2. E. J. Arthur, P. A. Hancock, and S. T. Chrysler. Spatial orientation in real and virtual worlds. In *Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting*, pages 328–332, 1993.
3. M. Pauline Baker and Christopher D. Wickens. Human factors in virtual environments for the visual analysis of scientific data. Technical Report NCSA-TR032, National Center for Supercomputing, August 1995.
4. Steffi Beckhaus, Felix Ritter, and Thomas Strothotte. Guided exploration with dynamic potential fields: the CubicalPath system. *Computer Graphics Forum*, 20(4):201–210, 2001.

5. Doug A. Bowman, David Koller, and Larry F. Hodges. Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. *Proceedings of the IEEE Conference on Virtual Reality*, pages 45–52, 1997.
6. William G. Chase. *Handbook of Perception and Human Performance, Vol II: Cognitive Processes and Performance*, chapter Visual Information Processing. John Wiley and Sons, 1986.
7. Luca Chittaro and Stefano Burigat. 3D location-pointing as a navigation aid in virtual environments. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*, pages 267–274, 2004.
8. Luca Chittaro, Roberto Ranon, and Lucio Ieronutti. Guiding visitors of Web3D worlds through automatically generated tours. In *Proceedings of the ACM Conference on 3D Web Technology*, pages 27–38, 2003.
9. Rudolph P. Darken and Barry Peterson. Spatial orientation, wayfinding, and representation. In Kay M. Stanney, editor, *Handbook of Virtual Environment Technology*. Lawrence Erlbaum, 2001.
10. Rudolph P. Darken and John L. Sibert. Wayfinding strategies and behaviors in large virtual worlds. In *Proceedings of the ACM CHI'96 Conference on Human Factors in Computing Systems*, pages 142–149, 1996.
11. Niklas Elmqvist, Mihail Eduard Tudoreanu, and Philip-pas Tsigas. Tour generation for exploration of 3D virtual environments. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, pages 207–210, 2007.
12. James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley, Reading MA, 1990.
13. Tinsley A. Galyean. Guided navigation of virtual environments. In *Proceedings of the ACM Symposium on Interactive 3D Graphics*, pages 103–104, 1995.
14. Andrew J. Hanson and Eric A. Wernert. Constrained 3D navigation with 2D controllers. In *Proceedings of the IEEE Conference on Visualization*, pages 176–182, 1997.
15. Lichan Hong, Shigeru Muraki, Arie Kaufman, Dirk Bartz, and Taosong He. Virtual voyage: interactive navigation in the human colon. In *Proceedings of the ACM SIGGRAPH Conference*, pages 27–34, 1997.
16. Takeo Igarashi, Rieko Kadobayashi, Kenji Mase, and Hidehiko Tanaka. Path drawing for 3D walkthrough. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 173–174, 1998.
17. Jock D. Mackinlay, Stuart K. Card, and Robertson Robertson. Rapid controlled movement through a virtual 3D workspace. In *Proceedings of the ACM SIGGRAPH Conference*, pages 171–176, 1990.
18. Randy Pausch, Dennis Proffitt, and George Williams. Quantifying immersion in virtual reality. In *Proceedings of the ACM SIGGRAPH Conference*, pages 13–18, 1997.
19. P. Peruch, J. L. Vercher, and G. M. Gauthier. Acquisition of spatial knowledge through exploration of simulated environments. *Ecological Psychology*, 7(1):1–20, 1995.
20. George G. Robertson, Jock D. Mackinlay, and Stuart K. Card. Cone trees: Animated 3D visualizations of hierarchical information. In *Proceedings of the ACM CHI'91 Conference on Human Factors in Computing Systems*, pages 189–194, 1991.
21. Roy A. Ruddle. The effect of trails on first-time and subsequent navigation in a virtual environment. In *Proceedings of the IEEE Conference on Virtual Reality*, pages 115–122, 2005.
22. Jürgen Schulze, Andrew Forsberg, Alexander Kleppe, Robert Zeleznik, and David H. Laidlaw. Characterizing the effect of level of immersion on a 3D marking task. In *Proceedings of HCI International*, pages 447–452, 2005.
23. Deyang Song and Michael Norman. Nonlinear interactive motion control techniques for virtual space navigation. In *Proceedings of the IEEE Virtual Reality Annual International Symposium*, pages 111–117, 1993.
24. Desney S. Tan, George G. Robertson, and Mary Czerwinski. Exploring 3D navigation: Combining speed-coupled flying with orbiting. In *Proceedings of the ACM CHI 2001 Conference on Human Factors in Computing Systems*, pages 418–425, 2001.
25. Edward C. Tolman. Cognitive maps in rats and men. *The Psychological Review*, 55(4):189–208, July 1948.
26. Colin Ware and Danny R. Jessome. Using the bat: A six-dimensional mouse for object placement. *IEEE Computer Graphics and Applications*, 8(6):65–70, 1988.
27. Colin Ware and Steven Osborne. Exploration and virtual camera control in virtual three dimensional environments. *Computer Graphics*, 24(2):175–183, 1990.
28. Eric A. Wernert and Andrew J. Hanson. A framework for assisted exploration with collaboration. In *Proceedings of the IEEE Conference on Visualization*, pages 241–248, 1999.
29. Bob G. Witmer and Michael J. Singer. Measuring presence in virtual environments: A presence questionnaire. *Presence*, 7(3):225–240, June 1998.
30. Ji Soo Yoon and Mary Lou Maher. A swarm algorithm for wayfinding in dynamic virtual worlds. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, pages 113–116, 2005.