Exploring Effects of Screen-Fixed and World-Fixed Annotation on Navigation in Virtual Reality

James Dominic*
School of Computing,
Clemson University,
Clemson, SC, USA

Andrew Robb†
School of Computing,
Clemson University,
Clemson, SC, USA

Abstract

In this paper, we consider the effect of different types of virtual annotations on performance during a navigation task in virtual reality. Two major types of annotations were shown to users: screen-fixed annotations that remained fixed in the user's field of view, and world-fixed annotations that are linked to specific locations in the world. We also considered three different levels of navigation information, including destination markers, maps visualizing the layout of the space being navigated, and path markers showing the optimal route to the destination. We ran a within-subjects study where participants completed three trials with each of the six combinations of annotation type and information level, for a total of 18 trials in a virtual environment. Average speed, distance traveled, and the time taken to reach the destination were recorded during each trial. Participants were also asked to point back to where they started the trial upon reaching the destination, as a measure of spatial memory. Finally, participants were tasked with completing a secondary activity while navigating, so as to assess what effect annotation types had on multitasking performance. Participants navigated significantly more quickly when using world-fixed annotations; however an interaction effect was observed between the type of annotation and the level of information, which suggests that world-fixed annotations are not inherently better than screen-fixed annotations; instead, it is important to consider both the type of annotation and what information it displays.

Index Terms: Human factors and ergonomics—locomotion and navigation—

1 Introduction

In virtual and augmented reality, annotations can be defined as additional information available to users that would not be present in an unaugmented experience. Annotations are common in virtual experiences, and are one of the most common uses of augmented reality (AR) [31]. For the purposes of this paper, we broadly classify annotations into two types: *screen-fixed annotations*, where the information is affixed to a position on the screen, and *world-fixed annotations*, where the information is affixed to an object or location in space. Each annotation type has its own advantages and disadvantages.

Screen-fixed annotations are presented on a display and are always visible to the user regardless of where they look in the world. Any information about an object or a task in the world can be displayed on a restricted fixed view plane using screen-fixed annotations. The area on the view plane where information is displayed is restricted to make sure that the annotations do not occlude the environment where the users perform the task. This is an effective method to make sure that the information is displayed to the user, however the amount of information which can be displayed is often limited. Displaying a lot of information in a restricted single view plane could lead to overcrowding of data [22]. *World-fixed annotations* are displayed overlayed or tagged to an object or location in the world and are visible to the user only when the user looks at the object which the annotation is tagged to. World-fixed annotations allow the developer to address the problem of overcrowding data at the cost of weakened legibility, as the information is no longer always available to the user.

These annotations can be better explained in the context of an AR application displaying price information of furniture, which uses a smart phone camera to look at various furniture in a room. Screen-fixed annotations can be used to display the price of the furniture on a corner of the phone screen. This works well when looking at one piece of furniture. However, while displaying price information of more than one piece of furniture using screen-fixed annotation, the user has to make relations between the prices shown and each piece of furniture in the view. This might leave the user confused and the application unsuccessful in conveying the price of each item in the room. World-fixed annotations allow the application to show the price of each furniture beside the furniture at the distance at which the furniture is shown. Thus making direct relations between price of each item and the item itself, at the cost of only showing prices for furniture that is currently in view.

In this paper, we consider how screen-fixed and world-fixed annotations affect performance in a spatial navigation task. Spatial navigation tasks are reasoning tasks where users must make decisions about how to move through the space around them [29]. When performing a spatial navigation task, screen-fixed annotations require users to map the information seen on the screen to the the space they are moving through. In contrast, world-fixed annotations already embed this information within the space itself. As such, we hypothesize that:

H1: Users will complete navigation tasks faster when supported by world-fixed annotations.

H2: Users will travel shorter distances during spatial navigation tasks when supported by annotations.

Performance in spatial navigation tasks are also highly influenced by the information made available to the user. As such, we also manipulate the amount of information available to the user. Based on common navigation paradigms, we consider three levels of information: when only the location of the destination is made available, when a map of the space is made available along with the location of the destination, and when the path to the destination to provided to the user. With regard to the amount of information available, we hypothesize that:

H3: Users will complete navigation tasks faster as they are provided with more refined levels of information.

H4: Users will travel shorter distances during spatial navigation tasks when provided with more refined levels of information.

*e-mail: domini4@clemson.edu
†e-mail: arobb@clemson.edu

2642-5254/20/$31.00 ©2020 IEEE
DOI 10.1109/VR46266.2020.00083

Screen-fixed annotations require users to divide their attention between the space they are moving through and the annotation. In contrast, world-fixed annotations embed the relevant information within the space, which may lead users to develop a stronger understanding of the space more quickly. As such, we hypothesize that:

H5: Users' spatial map of the space they have navigated will be stronger when supported by world-fixed annotations.

We hypothesize that world-fixed annotations will also increase performance on secondary tasks that must be completed while navigating to a destination. As world-fixed annotations keep users attention focused in the space, rather than on the screen-fixed navigation aids, we argue this may lead to better performance on these additional non-navigational tasks. As such, we hypothesize that:

H6: Users' will complete more secondary, non-navigational tasks when supported by world-fixed annotations.

2 Related Work

In this section, we discuss the background work done in screen-fixed annotations, world-fixed annotations, varying levels of information, and spatial understanding.

2.1 Annotations in VR and AR

Information Rich Virtual Environments (IRVEs) are demonstrated to better facilitate learning than traditional class room lectures [7, 18]. IRVEs are VEs which contain a lot of sensory information embedded into the environments. Annotations are virtual world extending elements, which add contextual information without the need of this information to otherwise fit to the augmented object in visual terms [15]. Annotations can help AR devices to bridge the gap between physical and virtual through metaphors such as "point and show" and "show and tell". The added-value of annotations is the ease with which digital information can be connected to objects in the world. We study two different types of annotations, which are screen-fixed annotations and world-fixed annotations.

Screen-fixed annotation is information about the world or an object in the world shown on the display (screen) and anchored on the display rather than on the object the information is about. Screen-fixed annotations appear as overlaid on top on the virtual world's projection [22]. These annotations do not give depth cues about the referent, but ensures visibility and legibility of the annotation. Bolton et al. found that highlighting landmarks using a screen-fixed annotations while driving improved response times and success rates by 43.1% and 26.2%, respectively, among drivers [5].

World-fixed annotations are those annotations which are tagged to a real or virtual object in the world (space) in AR or VR. In addition to helping users form spatial understanding with less cognitive effort in navigation tasks, world-fixed annotations help both users and applications to offload information to their surroundings, which otherwise has to be memorized by the participant or takes up fixed screen real-estate. World-fixed annotations have been applied to support doctors in several capacities, including overlaying 3D imaging data directly on top of the patient's anatomy [4], by highlighting the exact region of a patient's head needing to be operated on [28], and during multiple stages of wide varieties of vascular and oncologic intracranial pathologies [19]. Andersen et al. proposes a mentoring system for surgeons with the use of world-fixed annotations in AR [2]. The annotations where anchored to physical locations on the surgical site and would appear in an AR display in a tablet. The study found considerably lower placement error and focus shift while using the AR system with world-fixed annotations than using a fixed display without annotations.

There is substantial literature comparing the use of screen-fixed annotations and world-fixed annotations to display information in various fields like aircraft manufacturing by superimposing diagrams over real world objects [8], and in automotive navigation by projecting the navigation information spatially on the windscreen [12]. The navigation and other supplemental information is world-fixed on the windshield. Drivers exhibit a strong preference for systems which world-fixed the navigation information over systems which display the information on a fixed screen [14]. Bark et al. examined the effectiveness of see through 3D volumetric AR displays to provide navigation information for drivers [3]. The volumetric AR display displayed navigation information on top of objects in the world the user was navigating. Users were able to recognize turns earlier using the AR display compared to a center mounted screen based navigation system.

Merenda et al. evaluated the use of an AR display to provide navigation information to drivers while operating an automobile [21]. They compared the use of screen-fixed 2D display and a world-fixed AR display to designate a parking spot while driving. The users took same amount of time to spot the parking spot in both conditions. However, the user notable overestimated the distance to the parking spot while using the scree-fixed display compared to the world-fixed display. They conclude that world-fixed AR displays are more effective at conveying distance cues than screen-fixed displays.

The personal guidance system developed by Loomis et al. helps visually impaired individuals to navigate through familiar and unfamiliar environments without the assistance of guides [17]. The system uses audio signals to provide navigation information to the user. The system can provide non-spatialized audio instructions such as "straight", "left", or "right". The system could also provide spatialized audio instructions. If the system wants the user to turn left, the user will hear the utterance "one" coming from approximately 80 degrees from the left.Users navigated faster when navigation information was provided using spatialized audio compared to forms of conventional speech. Conventional speech has to be transformed by the user into the environment, whereas spatialized audio is overlaid on the environment.

While the above research suggests that world-fixed annotations are often more effective than screen-fixed annotations, work by Polys et al. illustrates how this may not always be true, particularly when dealing with large amounts of spatial data. Polys et al. found that displaying information on one layout space is better than tight spatial coupling of data in a VE [22]. Users performed search tasks in an IRVE which hosted an annotated cellular structures (e.g. nucleus, mitochondria, lysosomes). The users were able to perform search tasks faster and more accurately when the information was displayed in a single layout space, similar to screen-fixed annotations.

2.2 Levels of Information

In our study, we deal with three different levels of information which are presented to the users via world-fixed or screen-fixed annotations. Rijnsburger et al. [25] examines a system developed to provide personalized annotations for presenters during a presentation using a HMD. They had smaller annotations which were only 5 words or less and they had annotations which were 6 words or more. They also had annotations which showed only text, only images, or a text and image. The study found that users spent 36.5% time looking at annotations which has 6 words or more and a background image while they spend only 17% time for annotations up to 5 words. They also found that users preferred images over text and image and had an lower preference for text-only annotations. The study concludes that users spend more time looking at annotations and got more distracted as the annotations became more complex.

While navigating in a complex town center environment, users relied on landmark information most frequently [20]. Distance information and street names were infrequently used. The landmark information category included a wide range of information. The availability of information within the category was high. In addition, landmarks are traditionally used to provide navigation information in the real world. Distance information and street names were used to enable navigation decisions, but also to increase confidence and trust in the decision. Ruddle et al. found using a global map (which does not show smaller local objects) in conjunction with a local map (which shows local objects) to be the most effective way to navigate a very large scale VE [26]. The maps were equally effective while performing informed search. Users travelled 1/3 lesser distance while performing informed searches using global and local map together than they did while using either of the maps alone.

2.3 Spatial Understanding in Virtual Environments

Schnabel and Kvan [27] conducted a study with architects using a 2D representation and a 3D VE to understand and then reconstruct an architectural design. The study found that it is important for architects to use a tool which represents the three-dimensionality of the design they are working with. The reconstruction of the design was more complete for people who used the 2D representation of the space, however, the readability (or understanding) of the spacial volume was better done by participants who used the VE.

The type of navigation used in a VE will influence the user's understanding of the space. It has been found that use of techniques like virtual turns have shown a decrease in realism which could potentially result in reduced spatial orientation. Use of the human body for navigation in VR can result in better understanding of the virtual space. Riecke et al. [24] found that user experience significant gains in spatial perceptions when they are allowed to control rotations in virtual spaces with their bodies rather with joysticks.

Bowman et al. [6] conducted a study in which they examined different travel techniques used to navigate highly occluded VEs. They found that the use of a map before navigating a space resulted in poor performance in users when it came to spatial memory evaluation tasks. Though it seems counter-intuitive, the map could have given the users an illusion that they have an advantage and hence put less effort into understating the space, or the map itself could have been a source of cognitive load. They also found that in some cases the user only gave a cursory glance to the map and continued exploring the space without using the map.

3 METHODS

We designed a 2x3 within-subjects study, where users completed a spatial navigation task while support by different annotation types (screen-fixed vs. world-fixed) and with different levels of information: the location of the destination (destination), the floor plan (walls), and the path to the destination (path). Participants completed three trials for each condition, for a total of 18 trials. The order conditions were presented in was randomized for each participant. An example from each of the six conditions is shown in Figure 1a.

We display the location of the destination (destination) using the two different types of annotations. In the screen-fixed - destination (SF-D) condition (see Figure 1a), the destination was shown on a mini-map affixed the top left portion of the participant's screen. It was located about 1.5 meters in-front and about 0.5 meter above the head offset towards the left side of the participant. The minimap was overlayed and hence always displayed on top of the objects in the environment. The minimap had an indicator of the participants location and the direction they were facing. If the location of the destination was not visible on the screen, it was shown on the edge of the mini-map in the direction of its location. The world-fixed - destination (WF-D) condition (see Figure 1d) also only showed the destination, but this was visualized in the world by a marker placed at the location of the destination which was visible through walls.

The floor plan (walls) are also shown to the user using both types of annotations. In the screen-fixed - walls (SF-W) condition (see Figure 1b), the same map from the SF-D condition was shown, however the map now also included lines showing the walls and doorways in the area surrounding the participant. Similarly, the world-fixed - walls (WF-W) condition (see Figure 1e) extended the WF-D condition with a visualization of the floorplan of the space the participants were in. This was accomplished by highlighting the base of each wall using a shader that was drawn on top of all other objects in the space.

Finally, the path to the destination (path) was also shown to the user using both types of annotations. In the screen-fixed - path (SF-P) condition (see Figure 1c), the same map from the SF-D condition was shown, however a path was then drawn on the map to the destination (this path was calculated using the A* algorithm [10]). Similarly, the world-fixed - path (WF-P) condition (see Figure 1f) extended the WF-D condition with a visualization of the path to the destination drawn along the floor of the space. Unlike the visualization for the walls, the path was not visible through other objects.

3.1 Participants

Participants were recruited from a virtual reality class taught at [blinded]. A total of 48 participants (38 males) were recruited for the study, however 18 failed to complete the experiment due to simulator sickness (we consider this large dropout rate in our discussion). Thirty participants completed all 18 trials. Participants' were an average 22.7 ± 2.97 years old. Participants reported having spent an average of 4.7 ± 2.62 hours in VR prior to the experiment.

3.2 Apparatus

Participants interacted with the virtual environment using the HTC Vive Pro. The virtual environment participants navigated was created in the Unity game engine. We generated 18 different floorplans using the Dungeon Architect plugin available for Unity. The algorithm was tuned to produce complex blocks of smaller rooms connected by open hallways. Figure 2 shows one of the floorplans used in the study. Starting location and destination was also selected randomly for each trial, under the constraint that both the starting point and the destination must appear within one of the blocks of rooms, and they must be at least 72 meters apart. The average floorplan was 10,000 square meters.

Participants moved through the space using the arm swinging locomotion system provided by the Virtual Reality Toolkit 3.3.0 (VRTK) plugin [1]. In this locomotion method, participants move by swinging their arms while pressing the grip button on the HTC Vive controllers. This is a partial gait technique that mimics the motion our arms follow while normal walking. We selected this method for our experiment based on the results of Coomer et. al. who found that a similar locomotion method resulted in improved navigation performance and reduced sickness as compared to teleportation and joystick-based motion [9].

3.3 Procedure

Upon arriving at the study location, participants were introduced to the experiment and asked to sign an informed consent form. Upon signing the form, participants completed a brief demographics questionnaire and were asked to put on the HTC Vive. They were then placed in an open world without walls and were allowed to practice using the arm-swinging technique for motion throughout the space. Once they felt comfortable with it, they began the experiment.

Participants were told that, during the experiment, they would be taking on the role of a firefighter navigating through an unknown building that is on fire, with the primary goal of reaching the location of a trapped individual. As such, their primary task was to reach the target location as quickly as possible. However, as a secondary task, they were told to put out as many fires on the way as possible. This could be done using a held-held fire extinguisher, which would extinguish a fire after it had been sprayed for 1.5 seconds. The purpose of this secondary task was to evaluate the influence of the

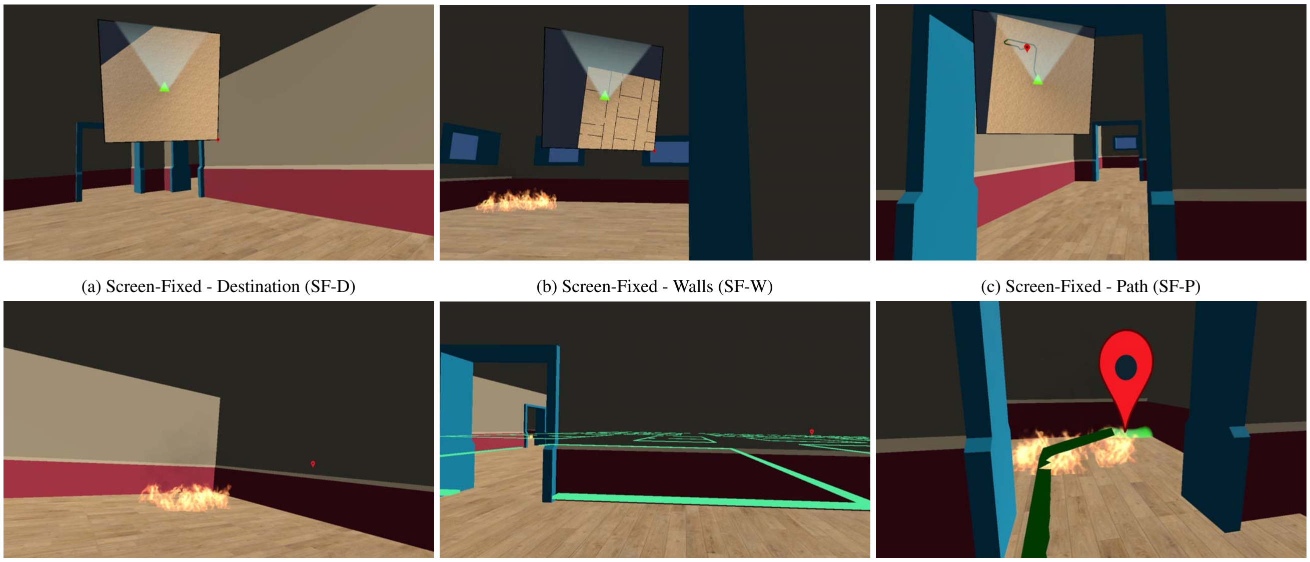

(a) Screen-Fixed - Destination (SF-D)

(b) Screen-Fixed - Walls (SF-W)

(c) Screen-Fixed - Path (SF-P)

(d) World-Fixed - Destination (WF-D)

(e) World-Fixed - Walls (WF-W)

(f) World-Fixed - Path (WF-P)

Figure 1: (a) shows destination information level presented through screen-fixed annotation. (b) shows destination walls level presented through screen-fixed annotation. (c) shows path information level presented through screen-fixed annotation. (d) shows destination information level presented through world-fixed annotation. (e) shows destination walls level presented through world-fixed annotation. (f) shows path information level presented through world-fixed annotation.

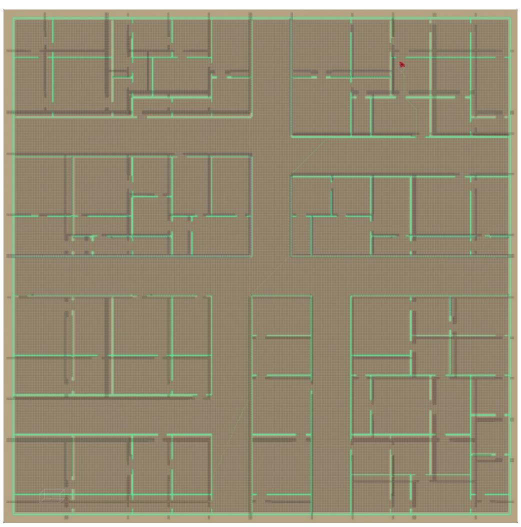

Figure 2: A sample floor plan deigned for the experiment

annotations on user performance on a secondary tasks while they were focused on the primary task, navigating through the space.

Finally, upon reaching their destination, participants were told to point towards their starting location. A ray was emitted from their controller to indicate the pointing direction, and their input was recorded with a button press on the controller. They then began the next trial. This continued until all 18 trials were completed, or until participants withdrew from the experiment due to sickness. Participants who completed all the trails took an average of 18 minutes and 45 seconds to complete the study.

Finally, participants completed a post questionnaire and a short debriefing. Participants who withdrew because of sickness were asked to complete these if they felt able to do so. All participants who withdrew because of sickness completed this portion.

3.4 Metrics

During each trial, the motion of participants through the virtual space was logged, along with the number of secondary tasks completed, the time taken to reach the destination, the total distance traveled, and the direction they pointed when indicating their starting location. We also logged the shortest possible path for that trial, based on the A* algorithm. We normalized the time taken to reach the destination, the number of secondary tasks completed, and the total distance traveled to account for the various shortest path lengths in different trials due to random position of the location the user started during each trial and variation in the floor plan size.

Participants completed a brief demographics questionnaire prior to the trials, and completed the IPQ presence questionnaire [23] and the Simulator Sickness Questionnaire (SSQ) [16] after the trials were completed.

4 Results

Linear mixed models were used to analyze the various data recorded while the participants navigated the complex floor plan. Annotation type and levels of information were used as fixed effects (including an interaction term). Participant ID was used as a random effect. To account for potential effects related to simulator sickness, we also included overall sickness as a fixed-effect in each of our analyses. Additional fixed effects were added to some models to control for additional factors relevant to the data being considered. These additional effects are discussed in the relevant sections. P-values were obtained by likelihood ratio tests of the full model against the model with the effect in question [30].

As the path participants followed was randomized during each trial, we normalized participants' data prior to analysis based on the length of the actual path they followed. Time taken and distance traveled were both normalized by dividing the participant's result by the length of the ideal path for that trial (as computed using the A* algorithm). As we expect time taken and distance traveled to be correlated with the actual length of the path, dividing by the distance of the ideal path removes the variable element between trials introduced by having paths of different lengths. The number of secondary tasks completed was normalized by dividing by the actual distance covered by participants, as the number of fires participants could encounter, and thus complete, is linked to the total distance traveled.

It should be noted that this resulted in some normalized distance traveled values that were lower than 1; this is an artifact of the limited granularity of the A* algorithm, which operates on a square grid and does not support diagonal paths. As such, it was possible (though rare) for participants to beat the "ideal" path length. That said, the result returned by the A* algorithm remained a reliable indicator of the relative path length across different trials, and thus remains suitable for use in normalizing the data.

4.1 Locomotion Behaviors

We selected time taken, distance traveled, and average speed as three measures to use to characterize participants' locomotion behavior, and how it was affected by the different annotations.

4.1.1 Time Taken

The time taken by each participant to reach the destination was recorded in each scene, and then normalized by the shortest possible path distance between the start and end points. The normalized number of secondary tasks completed was added as fixed effects to the model, to account for how completing more or less secondary tasks would influence the time taken to complete the experiment. Significant main effects were observed for both annotation type ($p < 0.001$) and level of information ($p < 0.001$). A mean of 0.425 and standard deviation of 0.179 was reported on normalized time taken to reach the destination. Participants took less time to navigate to the destination when using world-fixed annotations and as they received more refined levels of information. A significant interaction effect was also observed between annotation type and level of information ($p < 0.001$). Post-hoc pairwise comparisons revealed that participants performed significantly better in the WF-P condition than in the SF-D and SF-W conditions. A trend towards a significant effect was observed for simulator sickness ($p = 0.051$) on the time taken by the participant to reach the destination during the experiment, such that participants completed the task more slowly when experiencing higher levels of sickness. The time taken for each condition can be seeing in Figure 3a.

4.1.2 Distance Traveled

The distance traveled by each participant to reach the destination was recorded in each scene, and then normalized by dividing by the shortest possible path distance between the start and end points. The normalized number of secondary tasks completed was also added as a fixed effect to the model, to account for how completing more secondary tasks could add additional distance to the path participants traveled. A trend towards significance was observed with annotation type ($p = 0.057$) and a significant main effect was observed for level of information ($p < 0.001$). A mean of 1.154 and standard deviation of 0.389 was reported on normalized distance traveled to reach the destination. Participants traveled shorter distances as they received more refined levels of information, an there was a trend towards shorter distances when using world-fixed annotations. A significant interaction effect was also observed between annotation type and level of information ($p = 0.006$). Post-hoc pairwise comparisons revealed that participants performed significantly better in the WF-P condition than in the SF-D and SF-W conditions, but that the SF-P condition performed better than the WF-D condition. We observed no significant effect of simulator sickness ($p = 0.762$) on the distance traveled by the participant to reach the destination during the experiment. The distance traveled for each condition can be seeing in Figure 3b.

4.1.3 Average Speed

Average speed was computed by dividing time taken by actual distance traveled. The normalized number of secondary tasks completed was also added as a fixed effect to the model, to account for how completing more or less secondary tasks would influence the average speed of travel. Significant main effects were observed for both annotation type ($p < 0.001$) and level of information ($p < 0.001$). A mean of 3.107 and standard deviation of 1.103 was reported on average speed while navigating through the virtual space. Participants performed faster when using world-fixed annotations and as they received more refined levels of information. A significant interaction effect was also observed between annotation type and level of information ($p = 0.024$). Post-hoc pairwise comparisons revealed that participants moved more quickly in the WF-P condition than all other conditions. A significant effect of simulator sickness ($p = 0.034$) was observed on the speed at which participants traveled to the destination during the experiment, such that participants traveled more slowly when experiencing higher levels of sickness. The average speed for each condition can be seen in Figure 3c.

4.2 Spatial Memory

The error in participants' pointing towards their starting location was computed for each trial. The signed pointing error was recorded at the end of each trail, which was then converted to absolute pointing error for analysis. A trend towards significance was observed for annotation type ($p = 0.066$) but not for level of information ($p = 0.225$). A mean of 5.130 degrees and standard deviation of 45.307 degrees was reported on spatial memory task performance. The trend towards significance suggests that participants may have performed slightly better with the world-fixed annotations. No significant interaction effect was observed between annotation type and level of information ($p = 0.962$). However, we found a significant effect of simulator sickness ($p = 0.029$) on the spatial memory task performance of the participants during the experiment. The absolute pointing error for each condition can be seeing in Figure 4.

4.3 Secondary Tasks Completed

The number of secondary tasks completed by each participant while navigating the space was recorded in each scene, and then normalized by dividing by the distance traveled by the user. No significant main effects were observed for either annotation type ($p = 0.450$) or level of information ($p = 0.287$). A mean of 0.065 and standard deviation of 0.042 was reported on normalized number of secondary tasks performed while navigating through the virtual space. However, a significant interaction effect *was* observed between annotation type and level of information ($p = 0.040$), but post-hoc pairwise comparisons did not reveal any difference between specific conditions. An visual examination of the number of tasks completed, which can be seen in Figure 5, suggests that the interaction effect may have been caused by an increase in the number of tasks completed in the WF-P condition, as compared to all other conditions. We did not observe any significant effect of simulator sickness ($p = 0.922$) on the secondary tasks completion of participants during the experiment.

4.4 Survey Data

After participants had finished the trials, or withdrawn due to sickness, they completed the SSQ, a user preference survey, and the IPQ [23].

4.4.1 Simulator Sickness Questionnaire

Multiple one-way ANOVAs were run for each of the dimensions of the SSQ, and significant differences were found for each dimension

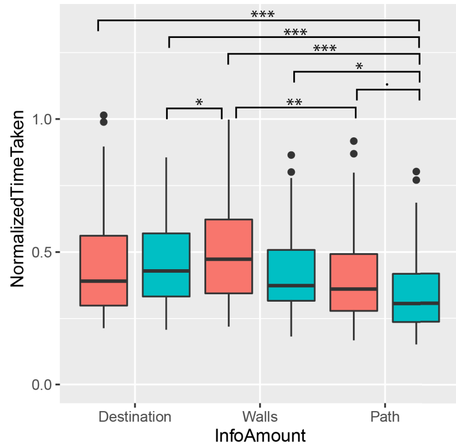

(a) The time taken by participants, divided by the shortest possible path distance for each trial, as determined by the A* algorithm, and then multiplied by 100 seconds.

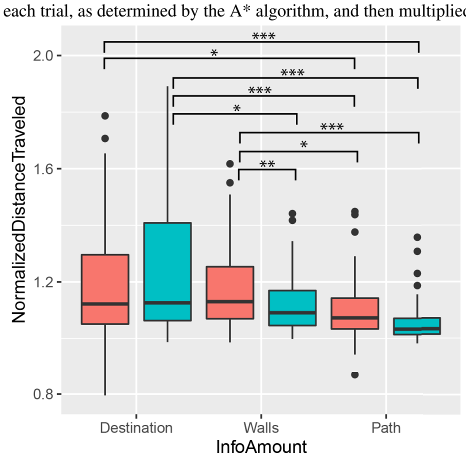

(b) Normalized distance traveled. A value of 1 represents the shortest possible path, as determined by the A* algorithm.

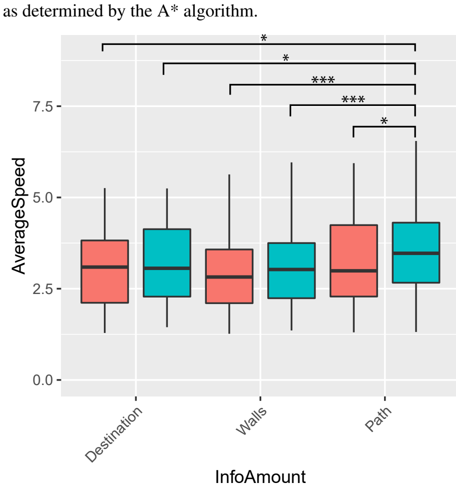

(c) The average speed of participants during a trials, as computed by dividing the time taken by the actual distance traveled.

Figure 3: These three metrics were used to characterize participants' locomotion behaviors and how they were affected by the different annotation types and information levels. Significant interactions are shown in the figures: * = p < 0.05; ** = p < 0.01; *** = p < 0.001

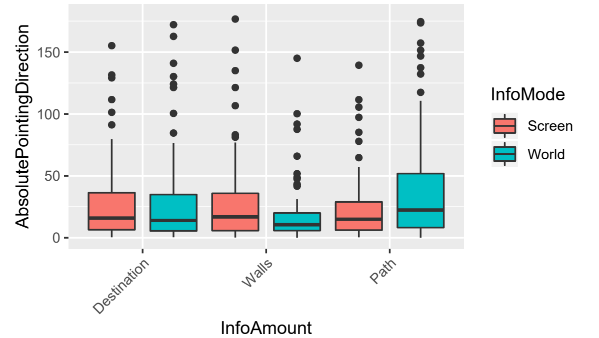

Figure 4: The deviation of direction pointed by the user from the direction the user actually started during each trial.

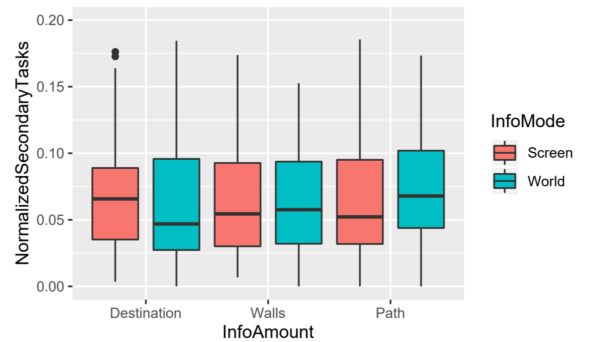

Figure 5: The number of secondary tasks completed by participants in each trial, normalized by the distance participants traveled.

($p_{Nausea} < 0.001$, $p_{Oculomotor} < 0.001$, $p_{Disorientation} < 0.001$, $p_{Total} < 0.001$). Participants who completed all of the trials reported low to acceptable levels of simulator sickness, while participants who withdrew early reported comparatively high levels of sickness. Figure 6 shows the simulator sickness scores recorded for participants who successfully completed and did not completed the experiment. The mean values of nausea, oculomotor discomfort, disorientation and total simulator sickness values reported by the participants are provided below.

Completed	Nausea	Oculomotor	Disorientation	Total
Yes	41.658	21.476	42.688	38.397
No	85.860	58.534	109.813	92.045

The standard deviation of nausea, oculomotor discomfort, disorientation and total simulator sickness values reported by the participants are provided below.

Completed	Nausea	Oculomotor	Disorientation	Total
Yes	39.333	25.342	59.275	41.819
No	37.736	36.473	61.494	45.501

4.4.2 User Preference Survey

Participants who successfully completed the experiment completed the user preference survey. Participants were asked to individually rate all navigation aids available to the them in the simulation on a scale from 1 to 5. A non-parametric Friedman test of differences among repeated measures revealed significant differences

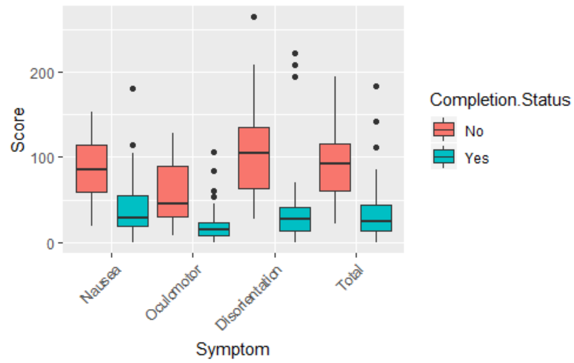

Figure 6: The simulator sickness values for participants who successfully completed and not completed the experiment.

between participants' ratings ($\chi^2 = 79.0, p < 0.001, W = 0.527$). Pairwise-comparisons using paired Wilcoxon signed-rank tests revealed significant differences between most conditions. However, no differences were observed between conditions with different annotations types but the same level of information (e.g. SF-P and WF-P). Additionally, no differences were seen between the WF-W and SF-P or WF-P conditions. This may suggest that user preferences were largely driven by the information displayed, rather than the manner in which it was displayed. Results from the user preference survey are shown in Figure 7.

5 Discussion

5.1 Performance on the Navigation Task

Our results show that both annotation type and information level impacted performance on the navigation task. Significant effect (or, in one case, a trend towards significance) was seen for both annotation type and information level on the time taken to complete the task, the distance traveled, and the average speed while traveling. A clear pattern was present in the data for annotation type: users performed better when using world-fixed annotations; they reached the destination more quickly, traveled shorter distances, and moved more quickly. These finding support H1 (*Users will complete navigation tasks faster when supported by world-fixed annotations*) and

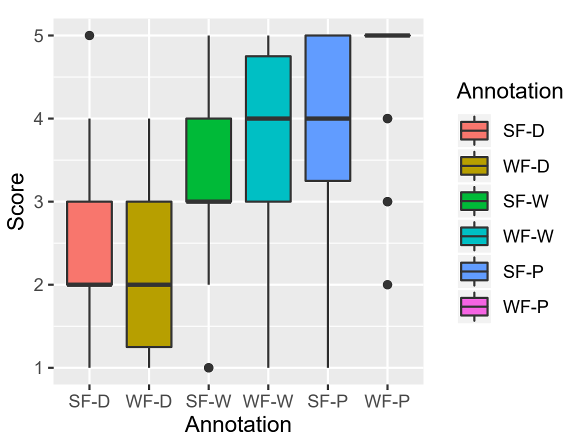

Figure 7: Box plot showing median and IQR of User preference for each type of navigation aid

H2 (*Users will travel shorter distances during spatial navigation tasks when supported by annotations*).

A pattern is less clear when considering the effect of information level. Information level had a significant impact on performance in all three of our metrics, however, it appears that the wall-based visualizations may have been outperformed by the destination-based visualizations on some metrics. That said, the path-based visualizations did outperform both destination-based and wall-based visualizations, which provides partial confirmation for H3 (*Users will complete navigation tasks faster as they are provided with more refined levels of information*) and H4 (*Users will travel shorter distances during spatial navigation tasks when provided with more refined levels of information*).

Several interaction effects were observed between annotation type and information level. In each of these interactions effects, the WF-P condition was shown to be significantly more effective than several other conditions for the given measure. WF-P was better than all other conditions for average speed, better than the SF-D and SF-W conditions for distance traveled, and better than the SF-D and SF-W conditions for time taken. This interaction effect suggests to nuance our observation that world-fixed annotations outperform screen-fixed annotations. Instead, a better interpretation is that *the right* world-fixed annotation will outperform screen-fixed annotations. Based on the interaction effects, the significant differences observed for annotation type appear to have been driven by the WF-P condition. The inverse of this interpretation also appears to be true: *the wrong* world-fixed annotation will underperform screen-fixed annotations. This is suggested by a significant interaction effect for distance traveled, where the a world-fixed condition (WF-D) is outperformed by a screen-fixed annotation (SF-P).

5.2 Performance on the Spatial Memory Task

A trend towards significance was observed for annotation type for the spatial memory task, suggesting that world-fixed annotations may have performed slightly better than screen-fixed annotations. However, without more supporting information, we cannot accept H5 (*Users' spatial map of the space they have navigated will be stronger when supported by world-fixed annotations*). Performance on this task was highly variable from trial to trial for each participant. We observed that participants were sometimes unsure where the starting location was, and would then point in a random location. Upon examining the data, we did not find any indication that specific participants ignored this task, as would be suggested by repeatedly exhibiting very poor performance. We hypothesize that this may have been caused by the design of the virtual levels: participants began navigation in a complex cluster of rooms, and ended navigation in a second complex cluster of rooms. It may be that some of these clusters required the participant to make such intricate navigational decisions that this hindered the spatial updating process, resulting in poor memory of their starting location.

It is interesting to note that performance was worst, on average, in the WF-P condition, which was the best condition for navigation performance. This could indicate that the low mental effort required by the WF-P condition resulted in poorer encoding of spatial information, which then hindered performance. The opposite is seen for the WF-W condition, which had the best average performance. Anecdotally, this condition imposed a higher cognitive load than most of the other conditions, which could have led to stronger encoding of the spatial information. However, without observing a significant interaction between annotation type and information level, this remains speculative.

5.3 Performance on the Secondary Task

No significant main effect of annotation type or information level was observed on the number of secondary tasks performed. However, a significant interaction effect *was* observed. Unfortunately, post-hoc tests did not revel a particular difference between conditions. Examining the data seen in Figure 5 suggests that this interaction effect may indicate that participants performed slightly more secondary tasks in the WF-P condition. If so this would partially support H6 (*Users' will complete more secondary, non-navigational tasks when supported by world-fixed annotations.*). However, the observed difference in performance is marginal at best. This potentially may be attributed to low overall completion rates for the secondary task. Approximately 25% of all trials had 0 secondary tasks completed. Users were instructed to put out fires while also moving to the destination as quickly as possible. Given that users were not penalized for failing to put out fires, it may be that participants chose to ignore this task in favor of the primary task.

5.4 User Preferences for Annotation Type and Information Level

In addition to performing the best, the WF-P was also the condition rated most highly by participants. A visual inspection of Figure 7 suggests that information level may have been the driving force behind the ratings given by users. Based on a rough grouping, the destination conditions were ranked the lowest, the wall conditions were ranked in the middle, and the path conditions were ranked the best. No similar delineation is clearly apparent for the world-fixed vs. screen-fixed conditions.

5.5 Limitations

Two limitations with this study need to be addressed. First, there was a high dropout rate due to simulator sickness (18 participants out of 48 participants); no data was used from participants who dropped out of the study in the primary analyses. Reported sickness was also included in our model to account for any variability it may have introduced. Our analysis of the reported sickness of participants who completed the study, as opposed to those who withdrew, showed that the participants included in our analysis reported significantly lower levels of sickness. In particular, their reported sickness was within the range of values commonly reported in VR studies, which suggests that our results will be relevant to any simulation that evokes typical levels of sickness.

As to why we experienced a high dropout rate due to sickness, this may be attributable to a combination of two factors: the use of screen-fixed annotations in combination with a partial-gait technique. Jerald et al. [13] and the Oculus best practices documentation [32] both report that using screen-fixed UI elements can lead to higher feelings of sickness in VR. In comparison to teleportation or real walking, partial gait techniques may lead to stronger feelings of sickness due to vection [11]. However, these techniques are commonly used in situations where it is desirable for users to have full, continuous control over their motion, but the tracked physical space is not large enough to support real walking. One user also commented that the perceived sense of speed in the environment contributed to his sickness: "I feel like I only felt discomfort because at some points I went really fast and didn't turn my head to run sideways".

Second, participants performed worse than anticipated on the spatial memory task and the secondary task. This makes it more difficult to assess the accuracy of H5 and H6. We believe that the complexity of the space participants navigated through may have contributed to the low performance on the spatial memory task, and that participants focus on the task of navigation resulted in low completion rate of the secondary task. It is also possible that the use of a partial gait locomotion technique could have contributed to the lower overall level of spatial awareness [24]. As such, special care should be taken to not interpret the lack of a significant result as an indication that no effect exists.

6 Conclusions

In this study, we evaluated the effect of world-fixed annotations and different levels of navigational information on performance in a navigation task, as well as on spatial understanding and the ability to complete secondary tasks while navigating. We found that users navigated more quickly and efficiently when using world-fixed annotations, and that the path-level of information yielded the best performance compared to the other levels of information. Perhaps most interestingly, we found interaction effects between annotation type and level of information that highlight how world-fixed annotations are not inherently better for navigation than screen-fixed annotations, but instead that the *right* world-fixed annotation can significantly improve performance while navigating, and that the *wrong* world-fixed annotation can also worsen performance. As such, the major contribution of this paper is to highlight the interaction that occurs between the information communicated by an annotation, and the means by which it is displayed to the user.

We did not find clear effects pertaining to the effect of annotation type and information level on spatial understanding or performance on secondary tasks. We discuss possible explanations for why no effects were observed here, and provide suggestions for future researchers interested in pursuing these questions further.

As VR and AR devices become more readily available, we expect to see them applied more often to support navigation in unfamiliar environments. While users may be unfamiliar with these environments, the systems they are using may not be, such as when a user visits a family member in a hospital, or when firefighters are working in a building that has already been mapped. Our results provide guidelines about how navigational information should be applied in these, and other similar, contexts. They can also be used to provide insights about what annotation types would be most appropriate for other levels of information, depending on the spatial components of that information.

References

[1] Vrtk - virtual reality toolkit. https://vrtoolkit.readme.io/.

[2] D. Andersen, V. Popescu, M. E. Cabrera, A. Shanghavi, G. Gomez, S. Marley, B. Mullis, and J. Wachs. Virtual annotations of the surgical field through an augmented reality transparent display. *The Visual Computer*, 32(11):1481–1498, 2016.

[3] K. Bark, C. Tran, K. Fujimura, and V. Ng-Thow-Hing. Personal navi: Benefits of an augmented reality navigational aid using a see-thru 3d volumetric hud. In *Proceedings of the 6th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, pp. 1–8. ACM, 2014.

[4] C. Bichlmeier, F. Wimmer, S. M. Heining, and N. Navab. Contextual anatomic mimesis hybrid in-situ visualization method for improving multi-sensory depth perception in medical augmented reality. In *2007 6th IEEE and ACM international symposium on mixed and augmented reality*, pp. 129–138. IEEE, 2007.

[5] A. Bolton, G. Burnett, and D. R. Large. An investigation of augmented reality presentations of landmark-based navigation using a head-up display. In *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, pp. 56–63. ACM, 2015.

[6] D. A. Bowman, E. T. Davis, L. F. Hodges, and A. N. Badre. Maintaining spatial orientation during travel in an immersive virtual environment. *Presence*, 8(6):618–631, 1999.

[7] D. A. Bowman, L. F. Hodges, D. Allison, and J. Wineman. The educational value of an information-rich virtual environment. *Presence: Teleoperators & Virtual Environments*, 8(3):317–331, 1999.

[8] T. P. Caudell and D. W. Mizell. Augmented reality: An application of heads-up display technology to manual manufacturing processes. In *Proceedings of the twenty-fifth Hawaii international conference on system sciences*, vol. 2, pp. 659–669. IEEE, 1992.

[9] N. Coomer, S. Bullard, W. Clinton, and B. Williams-Sanders. Evaluating the effects of four vr locomotion methods: joystick, arm-cycling,

point-tugging, and teleporting. In *Proceedings of the 15th ACM Symposium on Applied Perception*, p. 7. ACM, 2018.
[10] X. Cui and H. Shi. A*-based pathfinding in modern computer games. *International Journal of Computer Science and Network Security*, 11(1):125–130, 2011.
[11] J. L. Dorado and P. A. Figueroa. Methods to reduce cybersickness and enhance presence for in-place navigation techniques. In *2015 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 145–146. IEEE, 2015.
[12] A. Doshi, S. Y. Cheng, and M. M. Trivedi. A novel active heads-up display for driver assistance. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(1):85–93, 2008.
[13] J. Jerald. *The VR book: Human-centered design for virtual reality*. Morgan & Claypool, 2015.
[14] R. Jose, G. A. Lee, and M. Billinghurst. A comparative study of simulated augmented reality displays for vehicle navigation. In *Proceedings of the 28th Australian conference on computer-human interaction*, pp. 40–48. ACM, 2016.
[15] J. Keil, F. Schmitt, T. Engelke, H. Graf, and M. Olbrich. Augmented reality views: Discussing the utility of visual elements by mediation means in industrial ar from a design perspective. In *International Conference on Virtual, Augmented and Mixed Reality*, pp. 298–312. Springer, 2018.
[16] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology*, 3(3):203–220, 1993.
[17] J. M. Loomis, R. G. Golledge, and R. L. Klatzky. Navigation system for the blind: Auditory display modes and guidance. *Presence*, 7(2):193–203, 1998.
[18] G. Mantovani, R. Riva, and C. Galimberti. Vr learning: Potential and challenges for the use of 3d. *Towards cyberpsychology: Mind, cognitions, and society in the Internet age*, pp. 208–225, 2003.
[19] J. R. Mascitelli, L. Schlachter, A. G. Chartrain, H. Oemke, J. Gilligan, A. B. Costa, R. K. Shrivastava, and J. B. Bederson. Navigation-linked heads-up display in intracranial surgery: early experience. *Operative Neurosurgery*, 15(2):184–193, 2017.
[20] A. J. May, T. Ross, S. H. Bayer, and M. J. Tarkiainen. Pedestrian navigation aids: information requirements and design implications. *Personal and Ubiquitous Computing*, 7(6):331–338, 2003.
[21] C. Merenda, H. Kim, K. Tanous, J. L. Gabbard, B. Feichtl, T. Misu, and C. Suga. Augmented reality interface design approaches for goal-directed and stimulus-driven driving tasks. *IEEE transactions on visualization and computer graphics*, 24(11):2875–2885, 2018.
[22] N. F. Polys, S. Kim, and D. A. Bowman. Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. *Computer Animation and Virtual Worlds*, 18(1):19–38, 2007.
[23] H. Regenbrecht and T. Schubert. Real and illusory interactions enhance presence in virtual environments. *Presence: Teleoperators & Virtual Environments*, 11(4):425–434, 2002.
[24] B. E. Riecke, B. Bodenheimer, T. P. McNamara, B. Williams, P. Peng, and D. Feuereissen. Do we need to walk for effective virtual reality navigation? physical rotations alone may suffice. In *International Conference on Spatial Cognition*, pp. 234–247. Springer, 2010.
[25] W. Rijnsburger and S. Kratz. Personalized presentation annotations using optical hmds. *Multimedia Tools and Applications*, 76(4):5607–5629, 2017.
[26] R. A. Ruddle, S. J. Payne, and D. M. Jones. The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied*, 5(1):54, 1999.
[27] M. A. Schnabel and T. Kvan. Spatial understanding in immersive virtual environments. *International Journal of Architectural Computing*, 1(4):435–448, 2003.
[28] T. Sielhorst, M. Feuerstein, and N. Navab. Advanced medical displays: A literature review of augmented reality. *Journal of Display Technology*, 4(4):451–467, 2008.
[29] C. Ware and M. Plumlee. 3d geovisualization and the structure of visual space. In *Exploring Geovisualization*, pp. 567–576. Elsevier, 2005.
[30] B. Winter. Linear models and linear mixed effects models in r with linguistic applications. *arXiv preprint arXiv:1308.5499*, 2013.
[31] J. Wither, S. DiVerdi, and T. Höllerer. Annotation in outdoor augmented reality. *Computers & Graphics*, 33(6):679–689, 2009.
[32] R. Yao, T. Heath, A. Davies, T. Forsyth, N. Mitchell, and P. Hoberman. Oculus vr best practices guide. *Oculus VR*, 4, 2014.