

A Quantitative Assessment of the Impact on Spatial Understanding of Exploring a Complex Immersive Virtual Environment using Augmented Real Walking versus Flying

Victoria Interrante^{†1}, Eleanor O'Rourke², Leanne Gray³, Lee Anderson⁴ and Brian Ries¹

¹Department of Computer Science, University of Minnesota

²Department of Computer Science, Colby College

³Department of Computer Science, Kansas State University

⁴Department of Architecture, University of Minnesota

Abstract

When an immersive virtual environment spans an area that is larger than the available physical space for real walking, one may use an 'augmented walking' method such as Seven League Boots to enable participants to explore the space while gaining proprioceptive feedback that is similar to what they would experience with normal walking. In this paper, we present the results of a preliminary experiment in which we seek to quantitatively assess the extent to which participants are able to make more accurate spatial judgments about the locations of previously-seen targets in a complicated virtual city environment, experienced using a head-mounted display, after traveling to them using augmented real walking ('boots') versus virtual walking enabled by a button press on a hand-held wand. In a series of trials, we ask participants to follow paths of increasing complexity from a home base to different hidden targets in the environment and back. At each endpoint, with the path markings turned off, we ask participants to point, through the intervening alleyway walls, to the location they believe they started from. Participants are able to make real turns with their bodies in both locomotion conditions, however they are able to make real forward movement only under the augmented walking condition. Each participant completes eight trials under each locomotion condition, with the target locations and the order of experiencing each method counterbalanced between participants. In data collected from six participants so far, we are finding that the median angle error is significantly greater, overall, in the wand locomotion condition than in the 'boots' locomotion condition, and that the errors tend to increase, overall, as the path complexity increases (from two segments to four segments) in the wand locomotion condition but not in the 'boots' locomotion condition.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques

1. Introduction and Previous Work

Previous research has shown that participants tend to achieve a greater sense of presence in an HMD-based immersive virtual environment when they are enabled to travel through that environment using actual *real walking* than when they must make walking motions in place (*indirect walking*) or resort to pressing a button on a wand (*virtual walking*) [UAWB*99]. However, when the immersive virtual environment spans an area that is larger than the physically available navigable tracked space, using real walking becomes problematic.

Historically, the most common solution in those cases has been to either use an indirect technique for locomotion control, such as walking-in-place [TDS99] or leaning [LAKZ01], or to use a physical device such as an omnidirectional treadmill [e.g. DCC97]. Recently, several alternative approaches that seek to more faithfully

preserve the physical sensations associated with real walking have been proposed. The first is *re-directed walking* [RKW01], in which participants are surreptitiously induced to walk in a curving direction while believing that they are walking in a straight line. This method can be used to enable people to feel as if they are walking a long distance along a straight path, when they are actually just walking in a large circle. With the addition of large rotational distortion during large turns, this method can be successfully used to enable people to feel as if they are naturally walking between points in a space that is much more expansive than their actual traversed space.

The second approach is *augmented walking with seven league boots* [IRA07], in which participants experience accelerated movement through the virtual environment in the primary direction of their actual travel. In a prior preference study, we found that participants rated their experience of exploring a virtual hallway using seven

league boots significantly higher than their experience of exploring the hallway either using: ordinary (unscaled) walking but interrupted with turns when they ran out of space; virtual walking via a button press on a hand-held wand; or a more primitive form of augmented walking in which a uniform gain factor is applied to the output of the tracker (resulting in a sensation that one's head is wobbling exaggeratedly from side to side and up and down as one travels, in addition to moving forward more quickly). In this paper, we present the preliminary results of a follow-on study that seeks to quantitatively assess the extent to which enabling participants to explore an environment using seven league boots versus flying might enhance their ability to attain an accurate understanding of the spatial layout of the environment and their location within it.

2. Our Experiment

The general goal of our present experiment was to obtain quantitative measurements of participants' performance on a task requiring an intuitive spatial understanding of their environment, in order to seek objective insight into the question: under what conditions, and to what extent, might it be beneficial to enable users to move through a moderately large and complex immersive virtual environment using augmented walking with 'seven league boots' versus the standard default approach of enabling users to translate their viewpoint through such an environment using a hand-held wand?

2.1 Methods

The virtual environment used in this experiment consists of a collection of buildings of different sizes, laid out on a rectilinear grid with a complex maze of alleyways running between them (see figures 1a-c). Each building is assigned a random color and texture-mapped with a random brick pattern. Photographs of appropriately -sized and -placed windows and doors are superimposed over the walls of each building to provide participants with realistic, familiar cues to size in the virtual environment. The buildings are of different heights, all several stories tall, and they surround the participant on all sides at all times. A sky texture encloses the entire model to provide appropriate visual cues should a participant decide to look up (though no one has, yet).

Hidden throughout the environment, but made visible at appropriate times, are eight target items drawn from the SketchUp objects library. The target items include things like a lamp post, a bicycle, a dog house, an advertising kiosk, etc. Two *home base* locations, each denoted by a virtual patch of grass, are defined in the environment, on opposite sides of the city model. Four paths of varying complexity are defined from each home base to a target in the general vicinity of that location. Paths 1 and 2 are constructed so that they contain two straight segments joined by a single right angle turn; path 3 contains three straight segments, and path 4 contains four segments. The paths are denoted by colored stripes that extend down the middle of the alleyway, and they are made visible on an individual basis at appropriate times during each trial.

The experiment consists of eight trials, one for each pre-defined path. Participants complete four trials using

augmented real walking with seven league boots, and four trials using the wand to move forward. In all cases, participants must use their body to make a real turn as they follow a path around a corner; the main difference is that when they are using the boots to travel they must actually walk forward while pressing the trigger on the wand to engage the boots, while when they are only using the wand to travel they do not need to take any steps but just to press the trigger. The assignment of home base to navigation method is counterbalanced between participants, to control for any possible unintended effects of variations between the paths of equal complexity that might cause one to be somehow more difficult than the other. Also counterbalanced between participants is the order in which they experience each method. Explicitly, this means that: the first participant will begin by using the wand to follow each of the four paths from home base 1, and then use the boots method to follow each of the four paths from home base 2; the second participant will begin by using the wand to follow each of the four paths from home base 2, and finish by using boots to follow the paths from home base 1; the third participant will begin by using boots to follow each of the four paths from home base 1, and so forth.

At the start of each trial, the participant is positioned at a central point in our lab, which is registered with the appropriate home base in the virtual city model (we translate the location of the city with respect to the room between the two home base locations). One experimenter is assigned to manage the cables of the HMD and to watch that the participant doesn't walk into anything in the room. The other experimenter is seated at the computer and manages the transitions between the four stages in each trial. At the start of the first stage, the experimenter presses a button to make a path appear and the participant is asked to follow the path. Figure 1 shows what this looks like, from a point of view slightly behind the home base, which is just visible as a patch of grass at the bottom of the image.




Figure 1a: A screenshot showing what the participant sees at the start of a trial.

The path ends at the location of a target, which is made visible at the same time as the path leading to it. Figure 1b shows what this looks like, from a point of view displaced from the path so that both the turning of the path around the corner and the target can be seen in the image.




Figure 1b: A screenshot showing how one path leads the participant through the alleyways of the virtual city to a bicycle target

As the participant reaches the target, the path automatically turns off, and the participant is asked to turn around and orient himself or herself so that s/he is directly facing the location of the home base. At this point, the experimenter presses another button and a virtual wand is shown, at the position and orientation of the wand that the user is holding. S/he is asked to point the wand in the location of the home base and press the trigger. As the trigger is pressed, the wand turns red, so that the user can be certain that his or her response was recorded. The experimenter then presses another key to turn the path back on, and the participant is asked to follow the path back home, at which point the path again disappears. Once home, the participant is asked to again turn and point in the direction of the target that was just visited. This situation is shown in figure 1c.



Figure 1c: A screenshot showing what the participant sees at the stage in the experiment where s/he is asked to point toward the location where s/he believes the target to be. The small black line represents a pointer (modeled in 3D), which turns red when the trigger is pressed.

Finally, the participant is asked to close their eyes and travel, using their current locomotion mode, through the buildings to the location where they think the target is

located. The purpose of having our participants make this final judgment was to seek insight into the extent to which they were able to form an accurate mental model of the speed with which they were traveling through the environment using each navigation method.

Please note that the images 1a-c were taken from the auxiliary monitor that is used by the experimenters to keep track of what the participant is seeing while s/he is wearing the head mounted display. The slight barrel distortion visible in these images is used to compensate for the pincushion distortion caused by the optics in the HMD.

Figure 2 shows the setting within which the experiment was conducted. The virtual environment was presented via an nVisorSX head mounted display, which has an approximately 60° diagonal field of view and uses two LCOS displays of 1280x1024 resolution with 100% stereo overlap. Tracking was done with the HiBall 3000 from 3rd Tech, which provides high accuracy, low latency results. The maximum extent of the tracked space was approximately 30'x25'.



Figure 2: A photograph showing the setting within which the experiment was conducted.

Ten participants, two female and eight male, have completed the experiment so far, all of them undergraduate students recruited from a studio course in Architecture. Each participant was compensated with a \$10 gift certificate for his or her efforts.

3. Results

Figure 3, on the next page, illustrates the data we have collected so far from our experiment. Each colored dot represents the error in an individual positional estimate, and there is one dot for each judgment made by each participant. The horizontal bars amidst the dots represent the trimmed (80%) means of the absolute angle errors; these are displayed to provide an intuitive sense of the central error magnitude. Overall, the results plots suggest that participants appear slightly more likely to more accurately judge the direction to a place they have just come from when they are enabled to travel from one place to another in the virtual environment using augmented real walking than when they are restricted to virtually moving by pressing a button on a hand-held wand; however the effect is clearly not consistent or large. Furthermore, there is a huge apparent variability in the accuracy with which participants are able to make the type of directional

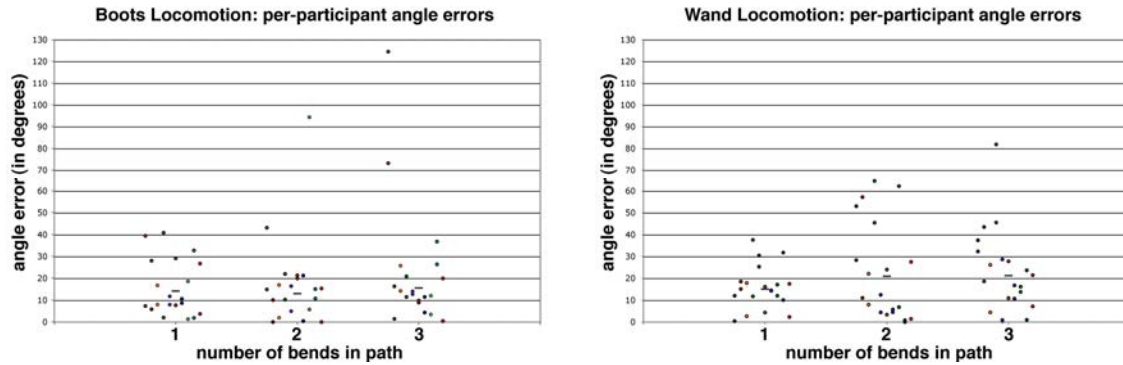


Figure 3: Plots of the errors in participants' estimates of the direction from a target to home base, and vice versa, as a function of the locomotion mode used to traverse the path between them and the path complexity. The horizontal bars in the midst of the dots represent the trimmed (80%) means of the errors in each case.

judgments requested, and we believe that this complicates our ability to rely on this metric to differentiate the experience in each case. In fact, recent research suggests that spatial cognition may rely on only very rough global Euclidean maps of space, if at all [RW06].

Nonetheless, to gain clearer insight into the potential significance of any effects of locomotion mode or path complexity on these spatial judgments, we performed a three-way analysis of variance (subject \times path complexity \times navigation mode) on our data after removing the most extreme outliers – individual direction estimates that differed in the magnitude of their error by over three standard deviations from the mean error computed over all other judgments made. We found a strongly significant difference in accuracy between participants $\{F(9,56)=2.83, p=0.008\}$, but only a marginally significant effect of locomotion mode overall $\{F(1,56)=3.08, p=0.085\}$, and no overall significant effect of path complexity $\{F(2,56)=0.95, p=0.392\}$.

Discussion and Future Work

Our finding that participants appear to be able to successfully achieve a reasonable, if coarse, spatial understanding of a virtual environment after exploring it along non-simple paths using augmented walking with seven league boots is consistent with similar findings by WNMC*06 who previously studied spatial cognition after augmented walking with uniform gain along simple straight paths in an open room VE.

Several promising directions for future work remain. Most importantly, we would like to more thoroughly explore the capabilities and limitations of augmented walking with seven league boots as a function of the amount of augmentation used. It seems reasonable to expect that the augmented walking metaphor might be most effective within a small range of scale factors closer to 1, to be completely ineffective when the amount of gain is too high, and to have intermediate effectiveness between these extremes, but this remains to be verified. Also, we would like to explore the use of alternative metrics for comparing the relative merits of augmented versus virtual walking.

Acknowledgments

This work was supported by a grant from the National

Science Foundation (IIS-0313226), and by the University of Minnesota through a Digital Technology Center seed grant and by the Linda and Ted Johnson Digital Design Consortium Endowment and Lab Setup Funds. Eleanor O'Rourke and Leanne Gray were supported in this work by the NSF-funded CRA-W Distributed Mentor Program, which enabled them to spend the summer of 2006 working at the University of Minnesota on this project.

References

- [CGBL98] CHANCE S.S., GAUNET F., BEALL A.C., LOOMIS J.M.: Locomotion Mode Affects the Updating of Objects Encountered During Travel: The Contribution of Vestibular and Proprioceptive Inputs to Path Integration. *Presence* 7(2): 168-178.
- [DCC97] DARKEN R., COCKAYNE W.R., CARMEIN D.: The Omni-Directional Treadmill: a locomotion device for virtual worlds. In *Proc. 1997 ACM User Interface Software and Technology*, pp. 213-221.
- [LAKZ01] LAVIOLA J.J., ACEVEDO FELIZ D., KEEFE D.F., ZELEZNIK R.C.: Hands-free multi-scale navigation in virtual environments. In *Proc. 2001 ACM/SIGGRAPH Symposium on Interactive 3D Graphics*, pp. 9-15.
- [IRA07] INTERRANTE V., RIES B., ANDERSON L.: Seven League Boots: A New Metaphor for Augmented Locomotion through Moderately Large Scale Immersive Virtual Environments. In *Proc. IEEE Symposium on 3D User Interfaces (2007)*, pp. 167-170.
- [RKW01] RAZZAQUE S., KOHN Z., WHITTON M.C.: Redirected Walking. In *Proc. Eurographics (2001)*, pp. 289-294.
- [RW06] ROTHMAN D.B., WARREN W.H.: Wormholes in virtual reality and the geometry of cognitive maps. In *Journal of Vision*, 6(6): 143a.
- [TDS99] TEMPLEMAN J.N., DENBROOK P.S., SIBERT L.E.: Virtual Locomotion: Walking in Place through Virtual Environments. *Presence* 8(6): 598-617.
- [UAWB*99] USOH M., ARTHUR K., WHITTON M.C., BASTOS R., STEED A., SLATER M., BROOKS, F.P. JR.: Walking > Walking-in-Place > Flying in Virtual Environments. In *Proc. ACM SIGGRAPH '99*, pp. 359-364.
- [WNMC*06] WILLIAMS B., NARASIMHAM G., MCNAMARA T.P., CARR T.H., RIESER J.J., BODENHEIMER B.: Updating orientation in large virtual environments using scaled translational gain. In *Proc. APGV (July 2006)*, pp. 21-28.