Ruddle, R. A., Volkova, E., & Bülthoff, H. H. (in press). Learning to walk in virtual reality. ACM Transactions on Applied Perception.

Learning to Walk in Virtual Reality

ROY A. RUDDLE[1,2], EKATERINA VOLKOVA[2], AND HEINRICH H. BÜLTHOFF[2,3]

Affiliations:

1 School of Computing, University of Leeds, UK

2 Max Planck Institute for Biological Cybernetics, Tübingen, Germany

3 Department of Brain and Cognitive Engineering, Korea University, Seoul, South Korea

---

This article provides longitudinal data for when participants learned to travel with a walking metaphor through virtual reality (VR) worlds, using interfaces that ranged from joystick-only, to linear and omnidirectional treadmills, and actual walking in VR. Three metrics were used: travel time, collisions (a measure of accuracy), and the speed profile. The time that participants required to reach asymptotic performance for traveling, and what that asymptote was, varied considerably between interfaces. In particular, when a world had tight turns (0.75 m corridors), participants who walked were more proficient than those who used a joystick to locomote and turned either physically or with a joystick, even after 10 minutes of training. The speed profile showed that this was caused by participants spending a notable percentage of the time stationary, irrespective of whether or not they frequently played computer games. The study shows how speed profiles can be used to help evaluate participants' proficiency with travel interfaces, highlights the need for training to be structured to addresses specific weaknesses in proficiency (e.g., start-stop movement), and for studies to measure and report that proficiency.

Categories and Subject Descriptors: I.3.6 [**Computer Graphics**]: Methodology and Techniques - *Interaction Techniques*. I.3.6 [**Computer Graphics**]: Three-Dimensional Graphics and Realism - *Virtual Reality*. H.5.2 [**Information Interfaces and Presentation**]: User Interfaces - *Input devices and strategies*.
General Terms: Experimentation, Human Factors, Performance
Additional Key Words and Phrases: Virtual reality interfaces, navigation, travel, metrics

---

This research was supported by an Alexander von Humboldt Fellowship for Experienced Researchers awarded to Ruddle, the Max Planck Society and the WCU (World Class University) programme through the National Research Foundation of Korea funded by the Ministry of Education, Science and Technology (R31-10008).
Authors' addresses: Roy A. Ruddle, School of Computing, University of Leeds, LS2 9JT, UK; Ekaterina Volkova, Max Planck Institute for Biological Cybernetics, Spemannstrasse 38, 72076 Tübingen, Germany; Heinrich H. Bülthoff, Department of Brain and Cognitive Engineering, Korea University, Seoul, South Korea.

1. INTRODUCTION

The need to navigate is intrinsic to virtual reality (VR) because, whether an application is for reviewing engineering designs, entertainment, social communication, planning or training [Blascovich and Bailenson 2011; Bowman et al. 2004; Stone 2002], if users are to accomplish the very purpose of using a given VR world then they have to view it from different places. A variety of metaphors for navigation has been proposed (e.g., walking, flying, scene-in-hand and eyeball-in-hand [Bowman, Kruijff, LaViola and Poupyrev 2004; Chen et al. 1988]), but most of the research into these metaphors has investigated their effect on the cognitive aspects of navigation (e.g., judge distances, remember a route, develop a cognitive map, or have an enhanced sense of presence) rather than their effect on users' ability to travel and maneuver.

Being able to travel while consuming minimal attentional resources indicates that a user is proficient at using a given interface, and lies at the heart of individual and gender differences in spatial knowledge acquisition that have been reported in studies of VR navigation [Waller 2000]. Studies that compare participants who regularly play first person shooter (FPS) computer games with those who don't, implicitly include travel proficiency as a factor in the study design [Smith and Du'Mont 2009], but there is a notable lack of research into the amount of time that users require to become proficient at traveling in VR worlds.

The present research aims to determine: (1) how users' proficiency at traveling changes over time, (2) how travel proficiency varies between interfaces, and (3) how travel proficiency should be assessed. The research's contributions are identifying fundamental differences between interfaces for the proficiency with which users travel, identifying metrics that characterize those differences and can be used to measure users' progress during training, suggesting training regimes that may address those differences, and highlighting the effect those differences may have had on the results of previously reported studies.

The scope of the present research is limited to travel that uses a walking metaphor, because that is the one that is most commonly used in navigation research. The following sections summarize previous research into metrics for assessing travel proficiency and the effect of different interfaces on that proficiency, and then report data from two studies in which participants underwent structured training to learn to travel through VR worlds before being asked to perform higher-level (route- and survey-knowledge) navigational tasks. The results of those tasks have previously been published {Ruddle, 2011 #820}{Ruddle, 2011 #757}, but the training data are new. The interfaces used in the studies ranged from joysticks, to real walking, and linear and omnidirectional treadmills.

2. TERMINOLOGY

Across the VR literature, different terms are sometimes used for the same types of interface. This paper adopts the following terminology for VR interfaces:

- *Desktop*: A user views a VR world on a computer monitor.

- *Immersive*: A user views a VR world using a display that largely excludes real-world visual cues and responds (at least in part) to the user's physical movements, e.g., a head-mounted display (HMD) or CAVE.
- *Travel*: Movement though a VR world using any interface.
- *View-direction travel*: A user can only travel in the direction in which they look (sometimes termed gaze-directed travel).
- *De-coupled travel*: A user may travel in a direction that is different to the one in which they look (e.g., strafing with a joystick, or where travel direction is defined by the orientation of the user's torso).
- *Joystick travel*: A user travels by moving a joystick forward, backward and sideways.
- *Actual walking in VR*: A user walks through an empty physical space that "contains" the VR world, so there is 1:1 correspondence between the user's movements in the real world and VR.
- *Walking-in-place*: A user makes a stepping motion to travel through a VR world, while remaining in one position in the real world.
- *Linear treadmill*: A conventional treadmill.
- *Omni-direction treadmill*: A treadmill on which a user can walk in any direction, so they physically both turn and translate.

3. RELATED WORK

Metrics for assessing travel proficiency come from three broad fields – VR itself, biomechanics, and human-computer interaction. In VR a number of standardized travel task tests have been proposed, with metrics based on time and accuracy [Bowman and Johnson 2001; Lampton et al. 1994]. Typically, time is measured for travel between two defined points or along a defined path. In some research travel was faster with a more sophisticated interface (one that de-coupled the travel and view directions, rather than only allowing view-direction travel) [Bowman and Hodges 1997], but in other research the opposite was true [Bowman and Johnson 2001]. A possible explanation is that participants in the latter study were given insufficient training, something that is shown explicitly in results from a spatial search study where participants who could travel forward, backward and sideways performed worse overall than participants who could only travel forward, but the difference between the types of travel was negligible once the first two trials had been completed [Lessels and Ruddle 2005].

Few studies have conducted a longitudinal investigation of different types of travel. One notable exception [Feasel et al. 2008], which compared walking in the real world with actual walking in VR, walking-in-place and joystick travel, highlighted differences in the amount of training required for different interfaces (least for actual walking) and the fact that in excess of 15 minutes of training was needed for walking-in-place and joystick travel. Of the 44 studies that were included in two reviews of VR navigation and involved active travel, as opposed to being passively transported, 30% did not provide participants with any practice in traveling through VR worlds before commencing the study itself [Ruddle 2011; Ruddle and Lessels 2006]. This is likely to have increased the variance of those studies' data, thereby reducing the likelihood of statistically significant differences being reported and/or biased the results in favor of natural/simpler interfaces. In fact, one study notes that participants who actually walked "needed the least time to familiarize themselves with the travel technique" and supports this with comments made by participants (e.g., "I never got used to the navigation!" for joystick travel) [Zanbaka et al. 2005].

When used as a VR travel metric, accuracy has typically been measured by counting the number of times a user collides with the world (e.g., a wall) or objects in it. These data may be either positively or negatively correlated with time, because collisions may slow down a user [Lampton, Knerr, Goldberg, Bliss, Moshell and Blau 1994], or the user may deliberately travel quickly and accept that collisions will occur (a speed-accuracy tradeoff). The latter is more likely if a VR system imposes no penalty for collisions, unlike the real world where obstacle avoidance is a requirement for people's survival [Pelah and Koenderink 2007].

Some studies have shown that fewer collisions occur with actual walking interfaces than desktop joystick travel or immersive view-direction travel with an head-mounted display (HMD) [Ruddle and Lessels 2009; Zanbaka, Lok, Babu, Ulinski and Hodges 2005]. However, other research that exclusively used an HMD showed no significant difference between actual walking and view-direction travel, but that more collisions occurred with participants who traveled where they pointed with their hand [Suma et al. 2010], indicating that it may be more appropriate to decouple users' view and travel directions by using torso rather than hand direction to dictate travel (e.g., see [Ruddle and Jones 2001]).

Biomechanics provides a variety of metrics that are based on displacement, velocity and acceleration, and concern either how a user travels (the dynamics of movement) or where (the path) [Winter 1990]. Amongst the metrics that have been applied to VR travel are stride length, speed and peak deceleration ("how"), and distance from a point or obstacle clearance ("where") [Feasel, Whitton and Wendt 2008; Fink et al. 2007; Whitton et al. 2005]. These metrics have mainly been used to analyze the extent to which VR interfaces that require the user to make a physical walking motion are identical to real-world walking. Only small differences are reported between actual walking in VR and walking in a real-world version of the environment [Fink, Foo and Warren 2007; Whitton, Cohn, Feasel, Zimmons, Razzaque, Poulton, McLeod and Brooks 2005]. In principle at least, a linear treadmill or walking-in-place interfaces could be tuned so that users' experience of VR travel closely resembles walking in the real world [Feasel, Whitton and Wendt 2008; Hollerbach 2002; Souman et al. 2010; Whitton, Cohn, Feasel, Zimmons, Razzaque, Poulton, McLeod and Brooks 2005], but travel using abstract devices (e.g., a joystick, mouse or keyboard) is clearly intrinsically different [Fink, Foo and Warren 2007].

Other metrics capture subjective aspects of travel. For example VR studies have asked participants to self-report the perceived ease of use of an interface, naturalness of movement, and presence [Bowman et al. 1997; Slater et al. 1995], and in mainstream human-computer interaction concepts such as "flow" are included in some usability questionnaires (e.g., "I felt in harmony with the environment"; [van Schaik and Ling 2005]. These metrics capture aspects of interface usage that are difficult to express in purely objective terms, and have been used to identify significant differences between redirection methods for walking-in-place [Peck et al. 2009], and that participants' sense of presence is significantly greater when they actually walk in VR than walk-in-place or fly [Usoh et al. 1999].

To summarize, two main points should be reiterated. First, although time and accuracy are widely used as metrics for travel, potential confounds (e.g., whether those metrics are negatively vs. positively correlated) mean that additional metrics are also required. Second, most previous research has only reported snapshots of VR travel (e.g., participants' average performance over a series of trials), so little is known about longitudinal changes in that performance. The remainder of this paper analyzes previously unreported training data from two studies. In the first each participant made 10 traversals of a 24 meter route to learn to travel [Ruddle, Volkova, Mohler and Bülthoff 2011], and in the second each participant made two traversals of a 270 meter route to learn to travel [Ruddle, Volkova and Bülthoff 2011]. Table I summarizes the interfaces that were used in each study, which were principally designed to investigate the effect of the translational vs. rotational components of body-based information on participants' route- and survey-knowledge. Joystick travel provides no body-based information, HMD view-direction travel provides the rotational component, a linear treadmill provides the translational component, and actual walking and an omni-directional treadmill provide both components. It follows that the HMD view-direction travel and linear treadmill interfaces were less immersive than the omni-directional treadmill interface, because of the components of physical body movement that were incorporated. Although both studies used the same general style of VR world (orthogonal virtual marketplaces), the first was more compact so the corridors along which participants traveled were substantially narrower (0.75m vs. 5m).

Interface	**Study 1**	**Study 2**
Joystick travel	Desktop display and gamepad joysticks to rotate (heading & pitch) and translate (forward, backward & sideways)	
HMD view-direction travel	Physically turn, wearing an HMD, but use a gamepad joystick to translate (forward, backward & sideways)	
Walking-based	Actual walking, wearing an HMD	Walk on a linear treadmill, wearing an HMD, but use a gamepad joystick to rotate (heading & pitch)
		Walk on an omni-directional treadmill, wearing an HMD

Table I. Travel interfaces used in the two studies.

4. STUDY 1: COMPACT WORLD

This study comprised two experiments that investigated the effects of landmarks (Experiment 1) and body-based sensory information (Experiment 2) on participants' ability to learn routes. The route-learning data have previously been published {Ruddle, 2011 #757}, but the interface training data reported here are new. Experiment 1 had four landmark conditions, but all participants used the same joystick travel interface and underwent identical training to learn to walk with that interface. Therefore, in the present paper, all the participants are termed as belonging to the same *Joystick-only* group. Experiment 2 investigated route learning under two conditions that both used an HMD: view-direction travel (*HMD-turn*) and actual walking (*HMD-walk*).

4.1 Method

4.1.1 Participants. A total of 107 people (50 men; 57 women) with a mean age of 25.7 years (*SD* = 5.7) took part in the study, but 12 participants withdrew from the study because of motion sickness. The data reported in this paper are for the other 95 participants. A total of 57 participants were in the Joystick-only group, 20 were in the HMD-turn group, and 18 were in the HMD-walk group. The data for route traversals 9 and 10 of one Joystick-only participant were lost due to a recording error. Participants were paid an honorarium for their participation. The study was approved by the local ethics committee.

4.1.2 Materials. Interior and plan views of the route that participants used to learn to walk are shown in Figure 1. The VR software was written using C++, ran on a Dell Inspiron M1710 laptop, and rendered the scene at 60 frames/second. Participants' position and orientation in the VR worlds was recorded at the same rate to a logfile for subsequent analysis.

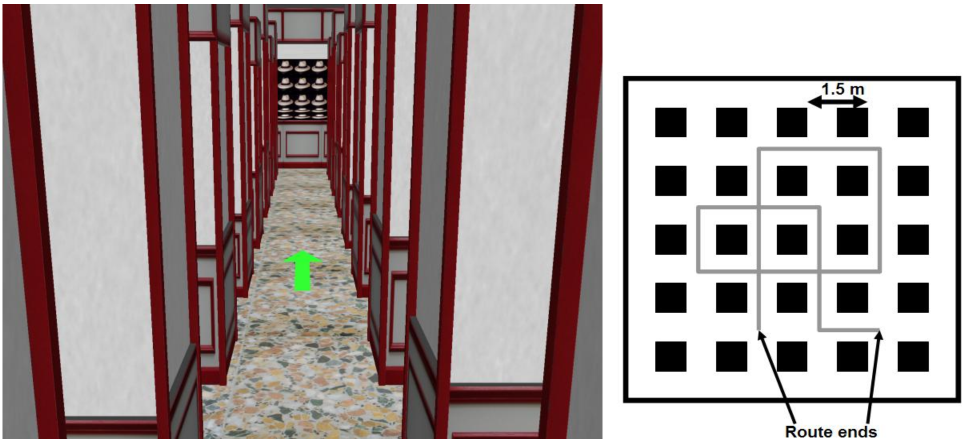

Fig. 1. Interior view (left) and plan view (right) of the 24 m route used in Study 1. Participants followed the arrow along the route, which crossed itself several times.

The Joystick-only group used the thumb-operated joysticks on a Logitech Rumblepad to control their movement. Manipulating the left joystick allowed participants to travel at a speed of up to 0.9 m/s (a slow walk) in any direction, and the right joystick allowed participants to vary the view heading and pitch. The heading and pitch could be changed at up to 120 and 25 degrees/second, respectively. Heading changes were seamless, but pitch was constrained so that participants could only look between vertically up and down. A non-stero 20-inch Dell flat panel display was used (1600 × 1200 pixels). The graphical field of view (FOV; 48º × 38º) was similar to the angle subtended by the monitor from a normal viewing distance (600mm).

The HMD-turn group stood in one place, viewed the VR world in stereo on an HMD, and moved by physically rotating (tracked by a Vicon MX13 motion capture system) and using one joystick on the Rumblepad to translate. The HMD was a nVisor SX (47º × 38º FOV; 100% binocular overlap; 1280 × 1024 pixels in each eye).

The HMD-walk group physically walked around a large tracking hall (see http://www.cyberneum.org) while viewing the VR world in the HMD. The position and orientation of a participant's head was tracked in six degrees of freedom using the Vicon system. For every group a slip collision response algorithm was implemented [Jacobson and Lewis 1997], so participants slid along objects if a collision took place rather than stopping instantaneously.

4.1.3 Procedure. First, the experimenter demonstrated how to traverse the interface practice route, using the desktop display and gamepad interface. Then participants traversed this route 10 times, first from A to B, then back to A, then back to B, and so on. The Joystick-only group always used that device and a desktop display. The HMD-turn and HMD-walk groups performed the first two traversals with the gamepad and desktop display, and the other eight traversals with an HMD and the interface they subsequently used in the main experiment.

4.2 Results

Three aspects of participants' travel were analyzed: time, accuracy and speed. To prevent any pause at the beginning of a trial from affecting the results, time was measured from when participants reached the first junction (1.5m from the start point) to the end of the route. Accuracy was measured by dividing the route into blocks (each block was from the center of one junction to the center of the next) and calculating the percentage of blocks in which a participant collided with the walls. This measured the amount of the route that participants had difficulty traveling along, whereas simply counting the number of collisions would not have differentiated between one participant making many collisions in a localized part of the route versus another participant making occasional collisions throughout the route. Speed was analyzed in terms of its profile and the percentage of time for which participants were "stationary" (traveling at 0.25 m/s or slower; tracking system noise means participants' would rarely have been absolutely stationary). The percentage collision and stationary data were normalized using an arcsin transformation prior to analyses of variance (ANOVAs) being performed. In the analyses below, only significant interactions are reported, Type III sum of squares was used because of the unequal group sizes, and a † after a *p* value indicates that the Greenhouse-Geisser correction was applied because the Mauchly sphericity test was significant. Games-Howell post-hocs were chosen because of the groups' unequal sizes and variance.

The longitudinal variation of the groups' mean times is shown in Figure 2. Data for the last eight traversals of the route were analyzed using an (ANOVA) that treated traversal as a within-participants factor and group as a between-participants factor. There were significant differences for group, $F(2, 91) = 5.87$, $p < .01$, traversal, $F(3, 307) = 19.61$, $p < .01^{\dagger}$, and a group × traversal interaction, $F(7, 307) = 2.26$, $p < .05$. Games-Howell post-hocs showed that the HMD-walk group took significantly less time than the HMD-turn ($p < .05$) and Joystick-only groups ($p < .01$).

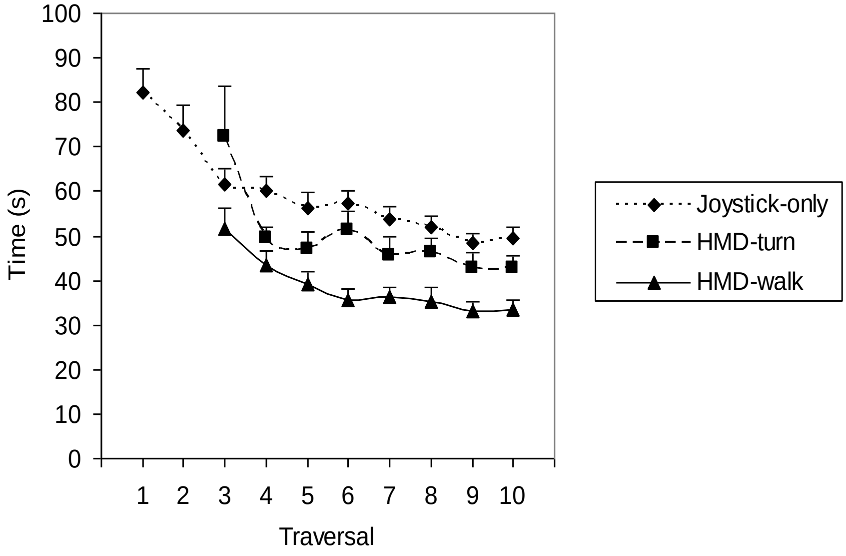

Fig. 2. Travel time for the traversals of Study 1. Error bars show standard error of the mean. Note: The HMD-turn and HMD-walk groups performed two traversals with the Joystick-only interface before performing eight traversals with their HMD interface.

The collision data (see Figure 3) was analyzed in the same way as the time. There was a significant difference for group, $F(2, 91) = 8.58$, $p < .01$, but not for traversal, $F(6, 12)$

= 1.56, $p > .05^{\dagger}$. Games-Howell post-hocs showed that the HMD-walk group collided with the walls in significantly fewer blocks on the route than the HMD-turn and Joystick-only groups ($p < .01$, in both cases). In previous studies, travel time and accuracy have sometimes been positively correlated and sometimes negatively correlated (see above). In the present study, an analysis of the Traversal 10 data for Joystick-only and HMD-turn participants (the HMD-walk group was excluded because only one participant in that group made any collisions) showed that travel time was negatively correlated with the percentage of the route with which participants collided, $r(75) = -0.20$, $p < .05$. In other words, there was a speed-accuracy trade-off.

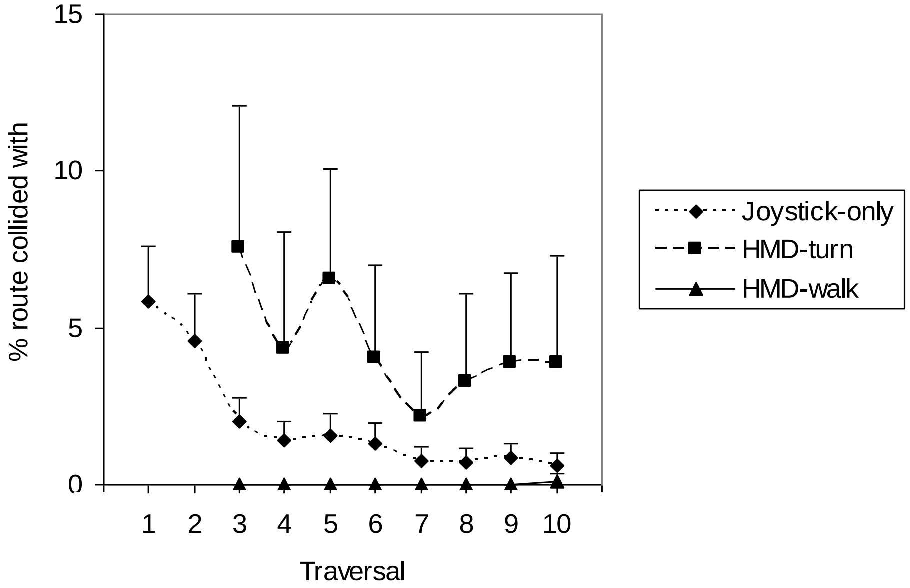

Fig. 3. Mean percentage of the route that participants collided with in Study 1, calculated from the normalized data and then untransformed. Error bars show standard error of the mean. Note: The HMD-turn and HMD-walk groups performed two traversals with the Joystick-only interface before performing eight traversals with their HMD interface.

To analyze how participants' speed varied, the percentage of time for each traversal that participants spent traveling in 0.25 m/s "bins" of speed was calculated (see Figure 4). This highlights a fundamental difference between the HMD-walk group and the other groups. The HMD-walk group spent only a small percentage of time stationary, even on the first walking traversal, and on the last traversal traveled at an average speed of 0.7 m/s. By contrast, the other groups spent almost half of the first traversal stationary, and even on the last traversal spent 31% (HMD-turn) and 35% (Joystick-only) of the time stationary. Analysis of changes in participants' view heading showed that the Joystick-only group did not turn for 43% of the time that they were stationary during the first traversal, reducing to 25% for the last traversal. Equivalent percentages for the HMD-turn were 44% and 26%, although it should be noted that this also included time when participants turned slowly (< 15 degrees/second), because sensor noise means that they are never completely still. When not stationary, the HMD-turn and Joystick-only groups moved at the maximum speed allowed (0.9 m/s; a slow walk).

The percentage of time that participants were stationary (i.e., speed ≤ 0.25 m/s) was analyzed in the same way as the time and collisions data. There were significant differences for group, $F(2, 91) = 71.12$, $p < .01$, and traversal, $F(5, 435) = 17.82$, $p < .01$[†]. Games-Howell post-hocs showed that the HMD-walk group was stationary for a significantly lower percentage of time than the HMD-turn and Joystick-only groups ($p < .01$, in both cases).

In a questionnaire, 11 of the participants reported that they played computer games frequently (at least once a week), and six of those were in the Joystick-only group and four were in the HMD-turn group. The difference in performance of these participants vs. the others (non-gamers) in those groups narrowed as the training progressed (see Table II). It is noticeable that even participants who played games frequently spent a substantial minority of the time stationary, even in the last four traversals.

Gaming	**Traversal time**		**% time stationary**		**% collisions**	
	Initial	**Last four**	**Initial**	**Last four**	**Initial**	**Last four**
Frequent	42 s	36 s	35%	29%	3%	2%
Not frequent	66 s	51 s	45%	36%	8%	5%

Table II. Traversal time, % time stationary and % of route collided with for participants in the Joystick-only and HMD-turn groups who played computer games frequently (at least once a week) vs. not frequently. *Initial* shows means for the first 6 (Joystick-only) or 4 (HMD-turn) traversals. For both groups, *Last four* shows means for the last 4 traversals.

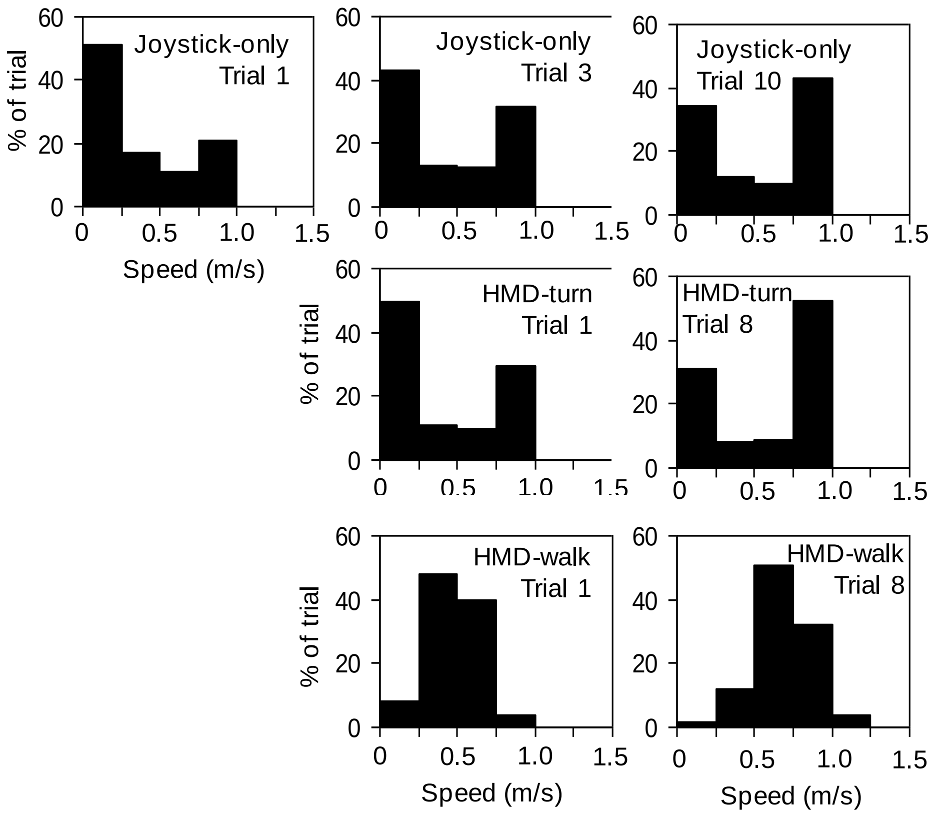

Fig. 4. Speed profile for each interface group in Study 1. Note: The HMD-turn and HMD-walk groups performed two trials with the Joystick-only interface before performing eight trials with their HMD interface.

4.3 Discussion

There was a similar pattern of results with all three metrics. The HMD-walk group traveled faster and made fewer collisions than the other groups, whose performance was equivalent to each other. The speed data showed that a key cause of the time difference was that the HMD-walk group moved smoothly along the route, whereas the other groups moved in a start-stop fashion and spent a substantial amount of time stationary.

After four traversals (3 minutes cumulative travel time) the HMD-walk group reached near-asymptotic performance, from then on spent negligible time stationary and, with one exception, never collided with the environment. These data support anecdotal evidence that, even without prior computer games experience, participants take immediately to actual walking for VR travel.

The Joystick-only and HMD-turn groups both exhibited difficulties traveling along the route, and even frequent computer game players spent one third of the time stationary. By the end of the training (7 minutes cumulative travel time) the time and collision data show that participants had reached an asymptote of performance that was inferior to the HMD-walk group. It could be that this asymptote difference reflects a fundamental difference between the way participants prefer to travel with those interfaces. Alternatively, the differences may have been narrowed if an alternative training regime was adopted, for example, instructing participants to move without stopping and turn while traveling, and only allowing them to 'pass' the training phase when they complete a non-stop traversal with no collisions.

Finally, the Joystick-only group used monocular viewing on a desktop display, whereas the other groups used stereo viewing in an HMD. The performance of the Joystick-only group relative to the HMD-turn group suggests that monocular viewing is not a disadvantage when this type of navigation is performed with desktop VR.

5. STUDY 2: BUILDING-SIZED WORLD

This study comprised two experiments that investigated the effect of body-based sensory information on participants' ability to develop a cognitive map. The cognitive map data have previously been published {Ruddle, 2011 #820}, but the interface training data from Experiment 2 that are reported here are new (Experiment 1 is not included because it did not use a training procedure with a prescribed route). There were four groups of participants, two with the same interfaces as the Joystick-only and HMD-turn groups described above, and groups that used a linear treadmill (HMD-linear) and an omnidirectional treadmill (HMD-omni; see Table I).

5.1 Method

5.1.1 Participants. Forty-four people (21 women) with a mean age of 26 years (*SD* = 5.1) took part, but four participants withdrew because of motion sickness. The data reported in this paper are for the other 40 participants (10 in each group). Participants were paid an honorarium for their participation. The study was approved by the local ethics committee.

5.1.2 Materials. Interior and plan views of the route that participants used to practice traveling are shown in Figure 5. The study used the same VR software as Study 1. For every group a slip collision response algorithm was implemented [Jacobson and Lewis 1997], so participants slid along objects if a collision took place rather than stopping instantaneously.

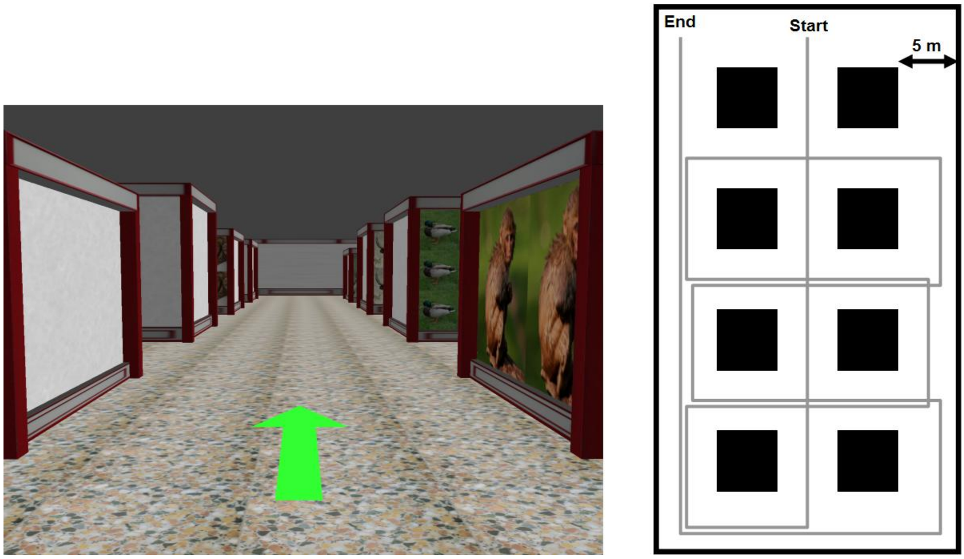

Fig. 5. Interior view (left) and plan view (right) of route used in Study 2. Participants traveled from the start to the end along a 270 m route that crossed itself several times.

The Joystick-only and HMD-turn groups' interfaces were identical to the one used in Study 1, except that the maximum speed was 1.34 m/s (faster than in Study 1), which was similar to the maximum speed of the treadmills. The HMD-linear participants walked on a 6 m long linear treadmill (see Figure 6a), which moved at participants' speed, and to look around or turn participants used the same device as the Joystick-only group. Guide ropes were used to help participants walk in a straight line along the treadmill. The HMD-omni group walked on a 4 × 4 m omni-directional treadmill (see Figure 6b) that moved at participants' speed, and participants were encouraged to walk normally. Both treadmills had control algorithms [De Luca et al. 2009; Souman, Giordano, Frissen, De Luca and Ernst 2010] that continually moved participants toward the center of the treadmill so they could walk at their own speed.

(a) (b)

Fig. 6. (a) The linear treadmill (HMD-linear), and (b) the Cyberwalk omni-directional treadmill (HMD-omni group).

5.1.3 Procedure. All participants started the experiment by practicing the experimental task (searching for objects in a virtual marketplace) using Joystick-only travel on a desktop display. This allowed the experimenter to explain the task face-to-face, and took an average of 11 minutes.

Next, participants practiced the traveling, using the interface for their group. For this, the HMD-omni group walked on the omni-directional treadmill with normal sight (no HMD) for 10 minutes, to get used to the way it operated, and then made two traversals of a defined 270 m route while wearing the HMD. The HMD-linear group walked on the linear treadmill with normal sight for two minutes to get used to the way it operated (less time was needed than for the omni-directional treadmill because walking on a linear treadmill is almost as straightforward as using one in a gym), and then made two traversals of a defined 270 m route while wearing the HMD. The HMD-turn and Joystick-only groups did not require any real-world familiarization (the former just had to turn, and the latter were seated) and, therefore, just practiced traveling through a VR world by making two traversals of the 270 m route using the interface and display (HMD vs. monitor) for their respective groups.

5.2 Results

As in the first study, participants' travel was analyzed in terms of time, accuracy and speed. Each analysis was a mixed factorial ANOVA that treated traversal as a within-participants factor and group as a between-participants factor. The data normalization, ANOVAs, and post-hocs were conducted in the same way as in Study 1 and, as before, only significant interactions are reported.

The mean travel time for every group in each trial is shown in Figure 7. There were significant differences for group, $F(3, 36) = 21.91$, $p < .001$, traversal, $F(1, 36) = 18.44$, $p < .001$, and a group × traversal interaction, $F(3, 36) = 5.73$, $p < .005$. Games-Howell post-hocs showed that the HMD-linear group took significantly more time than each of the other groups ($p < .005$ in every case).

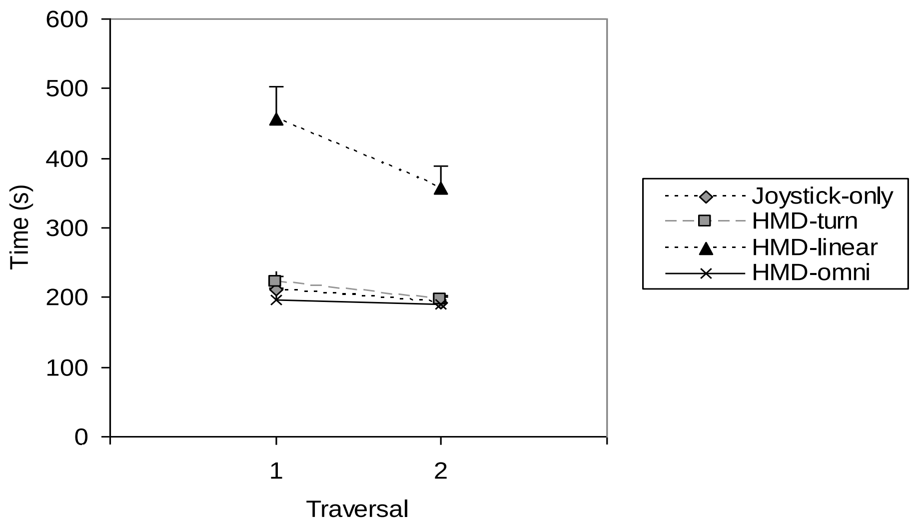

Fig. 7. Travel time for the traversals of Study 2. Error bars show standard error of the mean (NB: For all except the HMD-linear group, the standard error was small).

To further investigate whether groups had reached asymptotic performance in terms of travel time, the route was divided into two parts that had an equal length and number of turns (to do this, 10m in the middle of the route had to be included in both parts), and the time taken for the four half-routes (2 parts × 2 traversals) was compared (see Table III). This indicated that the HMD-linear group improved progressively from the first half-route to the last, whereas all the other groups had reached asymptotic performance before traversing the last half-route.

Interface	**Traversal 1**		**Traversal 2**	
	1st part	**2nd part**	**1st part**	**2nd part**
Joystick-only	117 %	102 %	101 %	100 %
HMD-turn	116 %	114 %	102 %	100 %
HMD-linear	147 %	120 %	108 %	100 %
HMD-omni	103 %	100 %	98 %	100 %

Table III. Traversal time in the first and second part of the route, expressed as a percentage of the time taken for the second part during Traversal 2. The percentages were calculated separately for each participant and then averaged for each group.

Participants' accuracy was measured in terms of the percentage of blocks in which a participant collided with the walls. There was no effect of group, $F(3, 36) = 2.03$, $p > .05$, or traversal, $F(1, 36) = 0.01$, $p > .05$. On average participants collided in 1% of the blocks, and 28 participants made no collisions at all, including all 10 participants in the HMD-linear group. The HMD-omni and Joystick-only groups contained eight and seven participants who made no collisions, respectively, but the HMD-turn group only contained three such participants.

Analysis of how participants' speed varied highlighted fundamental differences between the groups (see Figure 8). The Joystick-only and HMD-turn groups traveled at the maximum speed allowed by the joystick (1.34 m/s; an everyday walking pace) for the majority of each traversal. The HMD-linear group traveled slowly, but was rarely stationary, and the HMD-omni group showed the greatest variation in speed. An ANOVA showed that participants were stationary (i.e., speed ≤ 0.25 m/s) for a greater percentage of time during Traversal 1 than Traversal 2, $F(1, 36) = 5.61$, $p < .05$, but there was no effect of group, $F(3, 36) = 1.67$, $p > .05$.

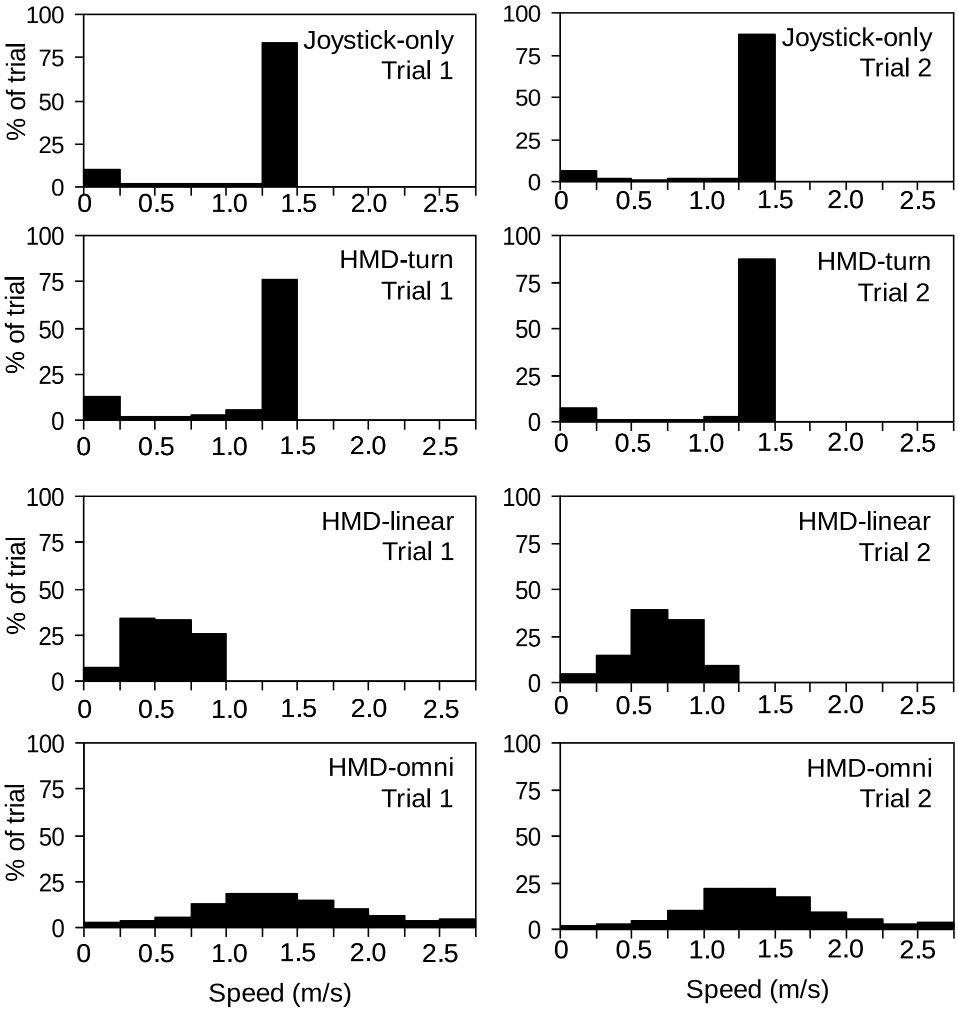

Fig. 8. Speed profile for each interface group in Study 2. For the HMD-omni group, the right hand bar shows the percentage of time spent traveling faster than 2.5 m/s.

5.3 Discussion

The travel time data highlight a stark contrast between the groups, with the HMD-linear group taking twice as long to travel along the route as the other groups did. The cause was due to the unnaturalness of the interface, which required participants to turn using a joystick even though they were viewing the VR world in an HMD, perhaps compounded by the guide ropes being a slightly elastic barrier to sideways movement. Rotational movements of a participant's head had no effect on the view that was rendered in the HMD, but translational movements did because they provided input to the treadmill control algorithm, which updated their position in the VR world. Even though the HMD-linear group improved with practice, the amount of time they spent learning (an average of 14 minutes over the 2 traversals) was not sufficient for them to travel as easily as participants in the other groups.

The other three groups took a similar time to complete the traversals, but the dynamics of their movement differed. The Joystick-only and HMD-turn groups mostly traveled at the maximum speed allowed by the joystick (1.34 m/s; similar to the maximum belt speed of the treadmills), although the number of HMD-turn participants who collided with stalls along the route is notable given that the width (5 m) of the corridors should have made maneuvering straightforward. The speed of the HMD-omni group varied more than any other group. High speeds are likely to have been caused by movements of a participant's head, when either looking around suddenly or compensating for movements of the treadmill (initially, some participants have a tendency to walk as if on a ship in rough seas), rather than their legs/body as a whole. The treadmill control algorithm takes as input a participant's position in the laboratory, measured by the Vicon system. Currently this position is calculated using markers on the HMD, but markers on a participant's waist would be better if problems caused by occlusion errors could be overcome.

6. GENERAL DISCUSSION

This paper analyzes training data from two studies of VR navigation, to determine how users' proficiency at traveling changes over time, how that proficiency varies between interfaces, and how travel proficiency should be assessed. The studies involved participants' navigation with interfaces that ranged from use of a joystick with a desktop display, to treadmills and actual walking while wearing an HMD. The environments all involved orthogonal arrangements of corridors, but varied considerably in terms of the frequency of turns and corridor narrowness (0.75 vs. 5 m). The present article effectively reports a field of study of navigation training, which adopted a hybrid approach (initial training using a Joystick-only desktop VR interface, and then progressing to group-specific interfaces) that our previous experience has proved to be pragmatic. As a result, the article analyses 'real' training for well over 100 participants, and relates some of the findings to previously published results about participants' performance in high-level (route and survey knowledge) navigation tasks. Limitations of our approach center on the fact that the results, therefore, do not show the effect of training when each interface is used in strict isolation of the others. However, even if the study had been purpose-designed for that, it is unlikely to have been practical to compensate for the different amounts of prior experience participants had with aspects of the setups that were used in the various conditions (e.g., desktop vs. immersive display).

The time that participants required to reach asymptotic performance for traveling, and what that asymptote was, varied considerably between interfaces. With an actual walking interface, participants on average took less than three minutes (4 traversals), even in a narrow environment that had had frequent, large (90°) turns. This illustrates the ease with which participants could maneuver in a confined space when using an interface that was "natural", and is consistent with the findings of previous research [Feasel, Whitton and Wendt 2008; Ruddle and Lessels 2009; Zanbaka, Lok, Babu, Ulinski and Hodges 2005].

With other interfaces, the time that participants required to reach asymptotic performance depended on the environment. When the corridors were wide then the quantity of training we provided (an average of 11 minutes with the Joystick-only interface (see §5.1.3) and 7 minutes with either the Joystick-only or HMDturn interface) was sufficient for most participants to travel collision-free at full speed. However, it was a different story when the corridors were narrow because, even after all of the interface training traversals had been completed (an average of 10 minutes/participant), participants were still taking substantially longer and making more collisions than participants who actually walked through the VR world (see Figures 2 & 3). This finding about training time is consistent with one of the few other longitudinal studies of VR travel training that has been published [Feasel, Whitton and Wendt 2008].

Difficulties maneuvering around sharp corners (90° turns in narrow corridors) with a mouse-based interface have been noted in previous research [Zhai et al. 1999] and, in the present study, even caused frequent gamers to spend a large amount of time stationary. This start-stop movement, which contrasts with the more continuous speed profile of the HMDwalk group (see Figure 4) is likely to inhibit path integration [Loomis et al. 1999] and reduce the extent to which participants can memorize a closely spaced sequence of turns as a single flowing “chunk” of movement. One can only speculate on the implications this holds for the results of previous VR studies, because the time taken to become proficient at maneuvering depends on the details of an interface and the world being navigated. However, given that the majority of studies provide less than 10 minutes of training (this was the case for at least 60% of the studies in the aforementioned reviews

[Ruddle 2011; Ruddle and Lessels 2006]), and the problems that some participants reported [Zanbaka, Lok, Babu, Ulinski and Hodges 2005], it is likely that many of those studies' results reflect participants' ability to develop spatial knowledge while they were still learning to travel, rather than reflecting the ability of participants who were fully trained in a given mode of travel [Waller 2000].

It follows that a lack of training may account for some of the significant differences reported in the test phase of the present studies. In Experiment 2 of [Ruddle, Volkova, Mohler and Bülthoff 2011], the HMDwalk group made significantly fewer errors than the HMD turn group and some of that difference may be attributed to a difference between the groups' interface proficiency. However, interface proficiency does not account for the marked advantage that the HMDwalk group showed on the return legs of the route, compared with the outward legs. It should also be noted that in between learning to travel (the data reported in the present paper) and the above test phase, participants made eight traversals of a 15 meter route to practice the type of task that was used during the test phase, and this also provided further practice at traveling. In Experiment 2 of [Ruddle, Volkova and Bülthoff 2011] there was a main effect of translational body-based information (HMDlinear & HMDomni groups vs. HMDturn & Joystick-only groups) for the accuracy of participants' straight line distance estimates, but it was also noted that the HMDlinear group's performance was suppressed and this was attributed to awkwardness of that interface. That is supported by data about the HMDlinear group's slow speed, which is reported in the present paper (see Figures 7 & 8).

Referring back to metrics that may be used to measure progress during training, the present study shows that the speed profile provides important insights into how participants travel that cannot be gained when just using time and/or accuracy metrics. The speed profile is different to the biomechanics-type metrics that have been used in previous research [Feasel, Whitton and Wendt 2008; Fink, Foo and Warren 2007; Whitton, Cohn, Feasel, Zimmons, Razzaque, Poulton, McLeod and Brooks 2005], and showed that in Study 1 the root cause of the time difference between the groups was the percentage of time for which the HMDturn and Joystick-only groups were stationary, whereas in Study 2 the root cause of the time difference was that the HMD-linear group simply traveled slowly.

The present article reports on participants' navigation with interfaces that ranged from use of a joystick with a desktop display, to treadmills and actual walking while wearing an HMD. The environments all involved orthogonal arrangements of corridors, but varied considerably in terms of the frequency of turns and corridor narrowness (0.75 vs. 5 m). To assess the extent to which the results can be generalized, we divide the interfaces into two categories: walking and non-walking. The walking interfaces (HMD-walk, HMD-linear & HMD-omni) all exhibited a Gaussian distribution for the speed profile, with participants' median speed affected by the specific interface (notably slow with the HMD-linear), but always increasing as training progressed. The profile reflected participants' ability to maneuver while traveling, which means that users should be able to navigate with similar ease in environments where they need to follow curved paths rather than just travel along straight line segments of corridor. We also predict that other types of walking interface such as walk-in-place {Feasel, 2008 #756} and redirected walking {Peck, 2009 #755} will have a similar speed profile. The non-walking interfaces were characterized by a bi-modal distribution (participants either traveled at full speed or were stationary), with the proportion of full-speed travel increasing with training and the width of the environment's corridors. Other studies have noted participants' tendency to travel in straight lines and inability to avoid obstacles when using non-walking interfaces {Zanbaka, 2005 #216}{Ruddle, 2009 #699}. Therefore, the problems that users have maneuvering seem to be inherent in interfaces that use abstract devices (e.g., a joystick, keyboard or mouse) to control translational movements, irrespective of whether or not an immersive display is used.

We conclude with some recommendations. The first is that studies should measure participants' travel proficiency using a blend of metrics that characterize both their speed profile and accuracy of travel. Second, studies should report these metrics to help readers judge the effect that travel proficiency may have had on the results of a study. Clearly it is more informative if longitudinal data are reported rather than a just an end-of-training snapshot. Third, training should be structured so that it directly addresses perceived weaknesses of participants' proficiency (e.g., start-stop movement), rather participants being left to their own devices. Fourth, studies should consider training participants to a given proficiency criterion, rather than for a fixed (and often nominal) length of time, to help compensate for differences between individuals' prior experience and proficiency.

ACKNOWLEDGMENTS

We thank Michael Kerger, Joachim Tesch and Betty Mohler for their technical assistance.

REFERENCES

BLASCOVICH, J. AND BAILENSON, J.N. 2011. *Infinite reality*. William Morrow, New York.

BOWMAN, D. AND HODGES, L. 1997. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *Proceedings of the 1997 Symposium on Interactive 3D Graphics* ACM, New York, 35-38.

BOWMAN, D. AND JOHNSON, D., LF 2001. Testbed evaluation of virtual environment interaction techniques. *Presence: Teleoperators and Virtual Environments 10*, 75-95.

BOWMAN, D., KOLLER, D. AND HODGES, L. 1997. Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. In *Proceedings of the IEEE Virtual Reality Annual International Symposium* IEEE, Los Alamitos, CA, 45-52.

BOWMAN, D.A., KRUIJFF, E., LAVIOLA, J.J. AND POUPYREV, I. 2004. *3D user interfaces: Theory and practice*. Addison-Wesley, London.

CHEN, M., MOUNTFORD, S., J AND SELLEN, A. 1988. A study in interactive 3-D rotation using 2-D control devices. *Computer Graphics 22*, 121-129.

DE LUCA, A., MATTONE, R., GIORDANO, P.R. AND BÜLTHOFF, H.H. 2009. Control design and experimental evaluation of the 2D CyberWalk platform. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2009)* IEEE, Los Alamitos: CA, 5051-5058.

FEASEL, J., WHITTON, M.C. AND WENDT, J.D. 2008. LLCM-WIP: Low-latency, continuous-motion walking-in-place. In *Proceedings of the 2008 IEEE Symposium on 3D User Interfaces* IEEE, Los Alamitos, CA, 97-104.

FINK, P.W., FOO, P.S. AND WARREN, W.H. 2007. Obstacle avoidance during walking in real and virtual environments. *ACM Transactions on Applied Perception 4*, article 2.

HOLLERBACH, J.M. 2002. Locomotion interfaces. In *Handbook of Virtual Environments: Design, Implementation, and Applications*, K.M. STANNEY Ed. Lawrence Erlbaum, Mahwah, NJ, 239-254.

JACOBSON, J. AND LEWIS, M. 1997. An experimental comparison of three methods for collision handling in virtual environments. In *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting* Human Factors Society, Santa Monica, CA, 1273-1277.

LAMPTON, D.R., KNERR, B.W., GOLDBERG, S.L., BLISS, J.P., MOSHELL, J.M. AND BLAU, B.S. 1994. The virtual environment performance assessment battery (VEPAB): Development and evaluation. *Presence: Teleoperators and Virtual Environments 3*, 145-157.

LESSELS, S. AND RUDDLE, R.A. 2005. Movement around real and virtual cluttered environments. *Presence: Teleoperators and Virtual Environments 14*, 580-596.

LOOMIS, J., KLATZKY, R.L., GOLLEDGE, R.G. AND PHILBECK, J.W. 1999. Human navigation by path integration. In *Wayfinding: Cognitive mapping and other spatial processes*, R. GOLLEDGE Ed. John Hopkins, Baltimore, MD, 125-151.

PECK, T.C., FUCHS, H. AND WHITTON, M.C. 2009. Evaluation of reorientation techniques and distractors for walking in large virtual environments. *IEEE Transactions on Visualization and Computer Graphics 15*, 383-394.

PELAH, A. AND KOENDERINK, J. 2007. Editorial: Walking in real and virtual environments. *ACM Transactions on Applied Perception 4*, 1-4.

RUDDLE, R.A. 2011. The effect of translational and rotational body-based information on navigation. *Manuscript submitted for publication*.

RUDDLE, R.A. AND JONES, D.M. 2001. Movement in cluttered virtual environments. *Presence: Teleoperators and Virtual Environments 10*, 511-524.

RUDDLE, R.A. AND LESSELS, S. 2006. Three levels of metric for evaluating wayfinding. *Presence: Teleoperators and Virtual Environments 15*, 637-654.

RUDDLE, R.A. AND LESSELS, S. 2009. The benefits of using a walking interface to navigate virtual environments. *ACM Transactions on Computer-Human Interaction 16*, article 5.

RUDDLE, R.A., VOLKOVA, E. AND BÜLTHOFF, H.H. 2011. Walking improves your cognitive map in environments that are large-scale and large in extent. *ACM Transactions on Computer-Human Interaction 18*, Article 10.

RUDDLE, R.A., VOLKOVA, E., MOHLER, B. AND BÜLTHOFF, H.H. 2011. The effect of landmark and body-based sensory information on route learning. *Memory and Cognition 39*, 686-699.

SLATER, M., USOH, M. AND STEED, A. 1995. Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Transactions on Computer-Human Interaction 2*, 201-219.

SMITH, S.P. AND DU'MONT, S. 2009. Measuring the effect of gaming experience on virtual environment navigation tasks. In *IEEE Symposium on 3D User Interfaces* IEEE, Los Alamitos, CA, 3-10.

SOUMAN, J.L., GIORDANO, P.R., FRISSEN, I., DE LUCA, A. AND ERNST, M.O. 2010. Making virtual walking real: Perceptual evaluation of a new treadmill control algorithm. *ACM Transactions on Applied Perception 7*, article 11.

STONE, R.J. 2002. Applications of virtual environments: An overview. In *Handbook of virtual environments*, K.M. STANNEY Ed. Lawrence Erlbaum, Mahwah, NJ, 827-756.

SUMA, E.A., FINKELSTEIN, S.L., REID, M., BABU, S.V., ULINSKI, A.C. AND HODGES, L.F. 2010. Evaluation of the cognitive effects of travel technique in complex real and virtual environments. *IEEE Transactions on Visualization and Computer Graphics 16*, 690-702.

USOH, M., ARTHUR, K., WHITTON, M., BASTOS, R., STEED, A., SLATER, M. AND BROOKS, F. 1999. Walking -> walking-in-place -> flying, in virtual environments. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* ACM, New York, 359-364.

VAN SCHAIK, P. AND LING, J. 2005. Five psychometric scales for online measurement of the quality of human-computer interaction in web sites. *International Journal of Human-Computer Interaction 18*, 309-322.

WALLER, D. 2000. Individual differences in spatial learning from computer-simulated environments. *Journal of Experimental Psychology: Applied 6*, 307-321.

WHITTON, M.C., COHN, J.V., FEASEL, J., ZIMMONS, P., RAZZAQUE, S., POULTON, S.J., MCLEOD, B. AND BROOKS, F.P. 2005. Comparing VE locomotion interfaces. In *Proceedings of the IEEE Virtual Reality Conference* IEEE, Los Alamitos, CA, 123-130.

WINTER, D.A. 1990. *Biomechanics and motor control of human movement*. Wiley-Interscience, New York.

ZANBAKA, C., LOK, B., BABU, S., ULINSKI, A. AND HODGES, L. 2005. Comparison of path visualizations and cognitive measures relative to travel techniques in a virtual environment. *IEEE Transactions on Visualization and Computer Graphics 11*, 694-705.

ZHAI, S., KANDOGAN, E., SMITH, B.A. AND SELKER, T. 1999. In search of the ‘magic carpet’: Design and experimentation of a bimanual 3D navigation interface. *Journal of Visual Languages & Computing 10*, 3-17.