Article:

Ruddle, RA, Volkova, E and Bulthoff, HH (2011) Walking improves your cognitive map in environments that are large-scale and large in extent. ACM Transactions on Computer - Human Interaction, 18 (2). ISSN 1073-0516

https://doi.org/10.1145/1970378.1970384

Roy A. Ruddle, Ekaterina Volkova, & Heinrich H. Bülthoff. (2011). Walking improves your cognitive map in environments that are large-scale and large in extent. ACM Transactions on Computer-Human Interaction, 18, 2, Article 10. http://doi.acm.org/10.1145/1970378.1970384.

Walking improves your cognitive map in environments that are large-scale and large in extent

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 DEPTARTMENT OF BRAIN AND COGNITIVE ENGINEERING, KOREA UNIVERSITY, SEOUL, SOUTH KOREA

CORRESPONDING AUTHORS:
ROY A. RUDDLE (R.A.RUDDLE@LEEDS.AC.UK),
AND HEINRICH H. BÜLTHOFF (HEINRICH.BUELTHOFF@TUEBINGEN.MPG.DE)

This study investigated the effect of body-based information (proprioception, etc.) when participants navigated large-scale virtual marketplaces that were either small (Experiment 1) or large in extent (Experiment 2). Extent refers to the size of an environment, whereas scale refers to whether people have to travel through an environment to see the detail necessary for navigation. Each participant was provided with full body-based information (walking through the virtual marketplaces in a large tracking hall or on an omni-directional treadmill), just the translational component of body-based information (walking on a linear treadmill, but turning with a joystick), just the rotational component (physically turning but using a joystick to translate) or no body-based information (joysticks to translate and rotate). In large and small environments translational body-based information significantly improved the accuracy of participants' cognitive maps, measured using estimates of direction and relative straight line distance but, on its own, rotational body-based information had no effect. In environments of small extent, full body-based information also improved participants' navigational performance. The experiments show that locomotion devices such as linear treadmills would bring substantial benefits to virtual environment applications where large spaces are navigated, and theories of human navigation need to reconsider the contribution made by body-based information, and distinguish between environmental scale and extent.

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, navigation, locomotion, cognitive map

1. INTRODUCTION

Most research into human navigation investigates our ability to perform tasks such as learning environmental layouts and routes in *large-scale spaces*. The defining characteristic of these spaces is that we have to travel through them to resolve the detail that is necessary for navigation, whereas a *small-scale space* has no visual barriers so all the navigational detail can be resolved from one position [Weatherford 1985] (Note: This is the definition of scale used in spatial cognition; in everyday life scale refers to a ratio between lengths). *Spatial extent* refers to an environment's physical area. Most environments that are large in extent are also large-scale (and vice versa), but exceptions are environments such as a car park or field (both large extent, but small-scale) and an open-plan office with tall partitions (small extent, but large-scale).

The distinction between scale and extent is particularly important when investigating the effect that body-based (proprioceptive & vestibular) sensory information has on navigation, because that information is used in path integration (the process of determining from your navigational movements how far and in which direction lays an earlier point on the path), and path integration errors increase with spatial extent [Loomis et al. 1999]. By contrast, the distinction between large- and small-scale navigation is purely visual, because once you can see how to travel somewhere the physical maneuvers involved are typically straightforward.

A long-standing, fundamental research question about three dimensional virtual environments (VEs; virtual reality worlds) is why are they so much more difficult to navigate than the real world [Lessels and Ruddle 2005; Suma et al. 2010; Witmer et al. 1996]? An important factor is the design of a VE's navigation interface, because that dictates the body-based information that is provided.

The present article describes two experiments that used different interfaces to investigate the effect of rotational vs. translational body-based information on participants' navigational performance (distance traveled) and cognitive map (direction and straight line distance estimates) when they searched for targets in a large-scale virtual marketplace. In Experiment 1, the marketplace was small in extent (9.75 × 6.75 m), so that it fitted within a tracking hall and one group of participants could literally walk around the virtual marketplace. Experiment 2 used novel treadmills, so that participants could navigate a virtual marketplace that was large in extent (65 × 45 m). No previous research has investigated the effect that the rotational and translational components of body-based information have on navigation in spaces that are large in both scale and extent. The results of the present study have profound implications for a wide range of VE applications, because navigation is an essential part of user interaction, and for our fundamental understanding of the sensory and cognitive processes that are involved in human navigation.

2. RELATED WORK

In navigation, a strong theoretical distinction is made between a person's knowledge of routes and their cognitive map [O'Keefe and Nadel 1978; Thorndyke and Hayes-Roth 1982]. Route knowledge is egocentric and, in its most basic form, is represented as a sequence of actions. However, the addition of landmark and metric information makes route knowledge more robust. A cognitive map (also termed survey knowledge, mental model, or mental map) represents an environment in an allocentric form, and provides information about the location of places within a world-reference frame. People are very good at learning real world environments (e.g., a new place of work, or holiday resort), because we develop a cognitive map from the outset [Montello 1998].

A cognitive map is particularly important for tasks such as exploring an environment (navigational search; [Ruddle and Lessels 2009]) and taking shortcuts [Foo et al. 2005]. During exploration, it is likely to be easier to remember where you have (not) traveled by forming a mental representation of the environment as a whole than remembering every single path segment you have traversed, and the efficiency of exploration may be quantified by calculating the distance traveled and the amount of repetition in a path. To take an effective shortcut, you need to set off in the correct direction and then travel an appropriate distance, at which point landmarks in the vicinity of your destination should be recognized. That is why two of the most common metrics for assessing the accuracy of a person's cognitive map are estimates of the direction to places and the straight line distance between places [Thorndyke and Hayes-Roth 1982].

The difficulty that people often have navigating VEs is assumed to be caused by a lack of environment fidelity and/or movement fidelity, when compared with the equivalent real-world setting [Waller et al. 1998]. The most important aspect of environment fidelity is the quantity of visual detail (landmarks, etc.) that is provided, and the effect of that on navigation is briefly reviewed in the following section. Movement fidelity relates to the design of a VE's navigation interface, and that dictates the body-based information that is provided. Previous research into the effect of body-based information on navigation is reviewed in more detail in §2.2.

2.1 Visual detail

The visual detail in an environment allows places to be identified and provides cues that could be used as landmarks. However, although landmarks do assist route learning [Jansen-Osmann and Fuchs 2006; Ruddle et al. in press], they provide much less benefit than is commonly assumed when the overall layout of a space needs to be learned [Ruddle and Lessels 2009; Ruddle et al. 1997]. Learning an environment's layout is a fundamental part of forming an accurate cognitive map.

A VE system's field of view (FOV) affects the quantity of a scene that may be seen at a given moment. When only visual information was provided, a wide FOV (180° horizontal) allowed participants to accurately perform a triangle completion task [Riecke et al. 2002] but, for a more complex navigational search task, a 144° FOV only led to a small improvement in performance when compared with a 48° FOV [Lessels and Ruddle 2004]. However, another difference was that the triangle completion study was performed using a physically large, projector display, but the navigational search study used three monitors. Other research has shown that participants remember the position of objects in a large-scale VE more accurately if the VE is viewed on a projected display rather than a monitor which subtends the same visual angle [Tan et al. 2006].

The virtual marketplaces used in the present study contained a rich variety of visual detail that could be used as landmarks, if participants wished. Due to the need for participants to physically navigate large virtual spaces, an head-mounted display (HMD) was used. This had a 47º × 38º FOV, which is typical for HMDs but less than can be achieved with curved projector displays.

2.2 Body-based information

The present article investigated the effect of rotational vs. translational body-based information on navigation. Desktop VEs provide almost no body-based information and, therefore, are typically termed *Visual-only*. By contrast, if the VE is viewed in HMD then, in some setups, the user physically turns but uses a joystick to translate. This provides body-based information for the rotational component of movement but none for the translational component, so the configuration is termed *Rotate*. Linear treadmills have, for many years, been advocated as a VE interface (e.g., [Brooks et al. 1992]), and may be used in conjunction with HMD or projector displays. A linear treadmill provides body-based information for the translational component of movement, but no body-based information for rotation, and so is termed *Translate*. Lastly, if a user physically walks through a VE while viewing it in an HMD then body-based information is provided for both components of movement, and this is termed *TransRot*. For Rotate configurations, the physical turning provides users with proprioceptive and vestibular cues. For TransRot configurations the cues depend on the movement interface. If users physically walk though a VE then they are provided with proprioceptive and vestibular cues for both the translational and the rotational component of body-based information. Walking-in-place removes the translational vestibular cues, and walking on a treadmill produces some conflicts in the translational vestibular cues (if, as in the present study, the treadmill operates at a user's speed then the vestibular cues will be correct when the user initially accelerates, conflicting while the treadmill adjusts its speed to gradually return the user to a central point on the treadmill, and then correct if the user maintains a constant speed). Like TransRot configurations, the cues provided by Translate configurations are interface-dependent. In the research that is summarized below, participants were provided with proprioceptive and vestibular cues for their configuration, unless otherwise noted.

Previous research into the effect of body-based information has used categories of environment that were:

a) small in scale and extent,

b) large-scale but small in extent, or

c) large in scale and extent.

In small-scale environments, influential research that used "optic flow" patterns as visual scenes suggests that the rotational component of body-based information is critical to prevent large, systematic errors from occurring during path integration [Avraamides et al. 2004; Klatzky et al. 1998]. However, studies conducted using rich visual scenes had markedly different findings, and suggest that translational body-based information is also required. In one such study, participants who walked around a virtual room (a TransRot group, in the terminology used in the present article) drew significantly more accurate sketch maps than participants who used a Rotate or Visual-only configuration [Zanbaka et al. 2005]. In another, where participants had to travel around a room to find targets in designated, possible locations, the TransRot group performed twice as many searches perfectly as Rotate and Visual-only groups [Ruddle and Lessels 2009].

With environments that were large-scale but small in extent (in the cases below, no larger than 15 × 13m), some research showed no difference between TransRot and Visual-only groups when participants had to remember the locations of objects after traveling a specific route or exploring a maze [Suma, Finkelstein, Reid, Babu, Ulinski and Hodges 2010]. However previous studies produced different findings, because participants' performance increased as the rotational and then translational component of body-based information was added. When participants traveled along a route and pointed to targets that had been encountered, a TransRot group pointed significantly more accurately than a Visual-only group, with performance of a Rotate group being in-between [Chance et al. 1998]. When participants had to learn a specific route after being guided along it once, a TransRot group made 36% fewer errors than a Rotate group, and behavioral data indicated that majority of the difference occurred because translational body-based information helped the TransRot group remember where to turn [Ruddle, Volkova, Mohler and Bülthoff in press].

In the real world, almost all large-scale environments are also large in extent, with examples being buildings, which are typically up to 100 × 100 m in size, villages (1 × 1 km) and cities (10 × 10 km, or greater). These are all at least one order of magnitude larger than the tracked laboratory spaces that have been used to study the role of body-based information in large-scale/small extent environments (see above), which limits the ecological validity of those studies. As extent increases, maneuverability becomes less important because obstacles are further apart, the time cost of making an error increases (this affects navigational behavior [Ruddle et al. 2000]), and there is greater opportunity for path integration errors to accumulate [Loomis, Klatzky, Golledge and Philbeck 1999].

Little previous research has investigated the effect of body-based information on the navigation of spaces that are both large-scale and large in extent. One study replicated some of the findings of [Chance, Gaunet, Beall and Loomis 1998], showing that a TransRot group of participants estimated directions to landmarks on a 840 m long route significantly more accurately than a Visual-only group [Waller et al. 2004]. However, this finding was confounded by the fact that the Visual-only group passively viewed scenes that were recorded during the navigation of the TransRot group, who actively navigated the route. In studies where all participants navigated actively, there was no significant difference in the distance that Rotate vs. Visual-only groups traveled to find target locations in virtual buildings and mazes, and no consistent difference between the groups' cognitive maps, as measured by estimates of direction and straight-line distance [Ruddle et al. 1999; Ruddle and Péruch 2004]. Similarly, when participants learned the layout of a virtual museum, there was no significant difference between the direction estimate accuracy of a Visual-only group and a TransRot group who walked-in-place (traveled by making a stepping motion while remaining in one place in the laboratory, meaning there was proprioceptive but little vestibular information for translational movement) [Grant and Magee 1998].

The following two experiments investigated the effect of rotational vs. translational body-based information on participants' navigational performance and cognitive map when they searched for targets in a virtual marketplace. In both experiments the marketplaces were large-scale environments, but the spatial extent was small in Experiment 1 and large in Experiment 2. The body-based information provided to each group of participants in the experiments is summarized in Table I.

Table I. Summary of the environments, movement interfaces, and body-based information used in the experiments.

<table>
<tr><th rowspan="2">Experiment</th><th colspan="2">Environment</th><th colspan="4">Body-based information group</th></tr>
<tr><th>Scale</th><th>Extent</th><th>Visual-only</th><th>Rotate</th><th>Translate</th><th>TransRot</th></tr>
<tr><td>1</td><td>Large</td><td>Small</td><td rowspan="2">Joystick & Desktop display</td><td rowspan="2">Joystick & HMD</td><td>-</td><td>Physically walk & HMD</td></tr>
<tr><td>2</td><td>Large</td><td>Large</td><td>Linear treadmill, joystick & HMD</td><td>Omni-directional treadmill & HMD</td></tr>
</table>

3. EXPERIMENT 1 (LARGE-SCALE; SMALL EXTENT)

A between-participants design with three groups was used (see Table I). The *TransRot* and *Rotate* groups viewed the virtual marketplace on a stereo helmet-mounted display (HMD) but, whereas the TransRot group had full body-based information (i.e., for translational and rotational movement), the Rotate group changed position using a joystick. The *Visual-only* group viewed the VE on a non-stereo monitor and changed position and orientation using two joysticks.

3.1 Method

3.1.1 Participants

Thirty-two individuals (10 women) with a mean age of 25 years (*SD* = 3.8) took part. All gave informed consent, took approximately 1½ hours to complete the experiment, and were paid an honorarium for their participation. The study was approved by the local ethics committee.

One Rotate participant (a man) withdrew because of motion sickness, and the data for a Visual-only participant (a woman) was discarded because she had great difficulty completing the task and traveled four times further than any other participant. The remaining participants were randomly assigned to each group, subject to the groups being gender balanced (7 men and 3 women in each).

3.1.2 Materials

The experiment took place in virtual marketplaces, which each comprised a grid of stalls, a long stall along each edge and four doors. The height of each stall was determined randomly (minimum = 2.0 m; maximum = 2.9 m), the length and width were both 0.75 m, and the corridor width was also 0.75 m. Every stall contained a picture of an everyday object that was visible from one side. The marketplaces were rendered at 60 frames/second using custom-designed software.

A marketplace with a 2 × 2 grid of stalls and one picture designated as a target was used to explain the task to participants. A 4 × 2 marketplace with two pictures designated as targets allowed participants to practice both moving around and the task. A 6 × 4 marketplace with four pictures designated as targets was used for the test (see Figure 1). Two versions of the test marketplace were constructed, with an identical layout but different pictures. Half of the participants in each group used each version.

Participants in the TransRot group physically walked around a large (13 × 12m) tracking hall while viewing the virtual marketplaces on an nVisor SX HMD (47º × 38º FOV; 100% binocular overlap; 1280 × 1024 pixels in each eye). The marketplaces were rendered by a Dell Inspiron M1710 laptop (NVIDIA GeForce Go 7950 GTX graphics card; Matrox DualHead2Go video splitter), which the experimenter carried in a backpack while walking behind each participant. Batteries powered the laptop and HMD, so the experimenter and participant traveled together as a wireless entity. The position and orientation of a participant's head was tracked using a Vicon MX13 motion capture system, and the participant's position/orientation in the VE was updated in real-time. Participants listened to white noise in headphones to mask any aural orientation cues from the hall, and were blindfolded when entering and leaving the hall so that they could not use knowledge of the general size of the hall to help memorize the environment.

Participants in the Rotate group listened to white noise, stood in one place and viewed the VE on the HMD. They traveled by physically rotating, which updated their orientation in the VE in real-time, but participants changed position using the left joystick on a Logitech Rumblepad (a common PC gaming device). The joystick allowed participants to travel at up to 0.9 m/s (a slow walk) in any direction.

Participants in the Visual-only group viewed the VE on a 20-inch Dell flat panel display (1600 × 1200 pixels), used the left joystick of the Rumblepad to change position and the right joystick to vary the view heading and pitch, at up to 120 and 25 degrees/second, respectively. The display was non-stereo and not head-tracked. The graphical FOV (48° × 38°) was similar to the angle subtended by the monitor from a normal viewing distance (600mm), and the HMD's FOV.

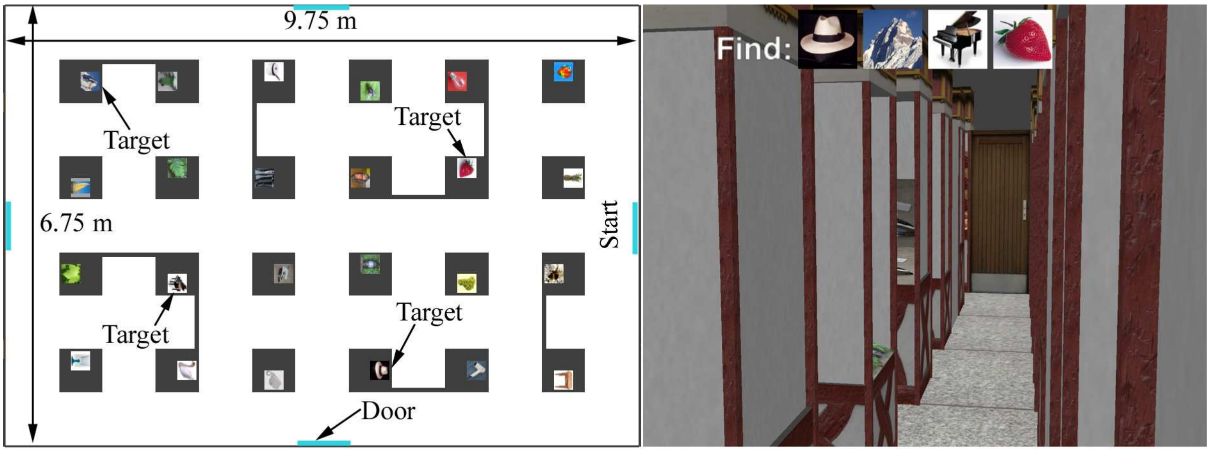

(a) (b)

Fig. 1. One of the test marketplaces: (a) Plan view (for illustrative purposes, the pictures have been made up-facing; they were actually placed as indicated in (b)), and (b) Interior view with pictures of the target objects visible at the top of the display.

3.1.3 Procedure

First, a participant performed two trials in the 2 × 2 marketplace. This was always done using the Visual-only configuration, so that the experimenter could explain the task face-to-face. Next, the participant practiced the task by performing two trials in the 4 × 2 marketplace, using the system configuration for their group (TransRot, Rotate or Visual-only). After this, the participant performed two test trials in one version of the 6 × 4 test marketplace and, finally, the participant answered a short questionnaire.

For each marketplace, in Trial 1 a participant searched for the target(s) in any order and then returned to the start point, indicating that they had arrived at each place by pressing a button on the Rumblepad. In Trial 2, the participant again searched for the target(s) in any order but, at each target, estimated the direction to every other target and the start point, and then estimated the straight line distance to the other target(s) and the start point. Once all the targets had been found, the participant returned to the start point and estimated the direction and straight line distance to each target.

Direction estimates were performed by either physically turning (TransRot and Rotate groups) or using the joystick to turn (Visual-only group) until the participant judged that they were looking through the stalls directly toward the specified target's location in the VE. In all cases, participants pressed a button on the Rumblepad to record their estimate. Distance estimates were reported verbally in meters and written down by the experimenter. There are many methods for recording and analyzing distance estimate data [Montello 1991]. The one used in the present study quantifies participants' knowledge of the relative straight line distances between places, which is known to be accurate when participants have well-developed cognitive map [Ruddle, Payne and Jones 1997; Thorndyke and Hayes-Roth 1982].

3.2 Results

Four types of data from the test marketplaces were analyzed: (a) the distance that participants traveled, (b) the accuracy of participants' direction estimates, (c) participants' sense of relative straight line distance, and (d) the speed at which participants traveled. The data types (b) and (c) are widely used as a measure of the accuracy of participants' cognitive maps [Thorndyke and Hayes-Roth 1982]. There are a variety of ways that participants' knowledge of distances may be assessed [Montello 1991]. Verbal estimates of absolute distance are error prone even when a participant has an accurate cognitive map, but the correlation of those estimates with the actual distances produces a pattern of results that is consistent with other measures of spatial learning [Thorndyke and Hayes-Roth 1982]. In the present experiment, all the data were analyzed using analyses of variance (ANOVAs) and there were no significant interactions.

The distribution of the distance traveled data was normalized using a natural logarithm transformation. A 3 × 2 (group × trial) mixed factorial ANOVA showed main effects of group, $F(2, 27) = 5.31$, $MSE = 0.20$, $p = .01$, $\eta_p^2 = .28$, and trial, $F(1, 27) = 29.82$, $MSE = 0.07$, $p < .001$, $\eta_p^2 = .52$ (see Figure 2). Tukey HSD post-hocs showed that the TransRot group traveled significantly less distance than the Rotate ($p = .005$) and Visual-only groups ($p = .02$), these latter two groups being statistically equivalent. All three groups traveled less distance in the second trial, but the TransRot group outperformed the other groups in both trials. Only two participants visited the targets in an identical order in both trials, and none reversed the order. Transformed back from logarithms and expressed as multiples of the shortest path, the distances for Trials 1 and 2 were 2.2 and 1.4 (TransRot), 3.4 and 2.3 (Rotate), and 2.9 and 2.2 (Visual-only).

The questionnaire showed that four participants played computer games frequently (at least once a week). In terms of total distance traveled they were ranked (with 1st being the participant who traveled the least distance) as follows within their group (overall): TransRot group 2nd (3rd) and 3rd (4th), Rotate group 7th (26th), and Visual-only group 3rd (11th).

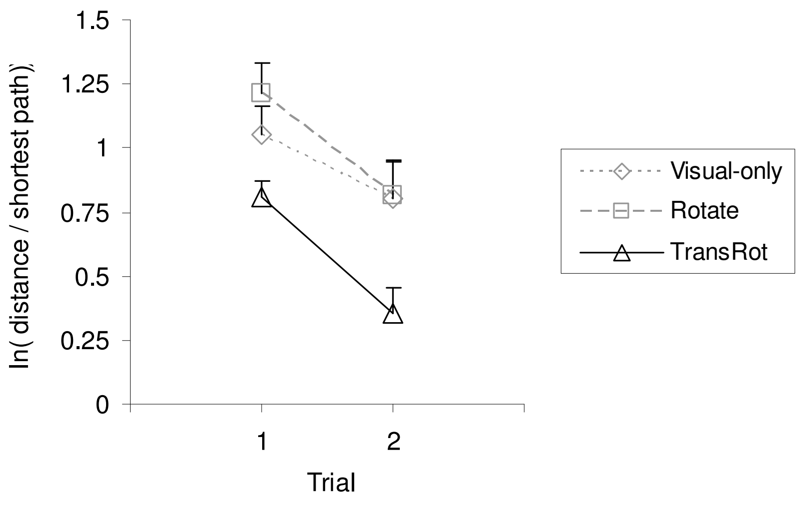

Fig. 2. Natural logarithm of the distance participants traveled during Experiment 1, expressed as multiples of the shortest path (for the shortest path, ln(distance) = 0.0). Error bars show standard error of the mean.

From participants' direction estimates, the mean absolute angular error was calculated. The distribution of these data was normalized using a natural logarithm transformation, and a univariate ANOVA showed a marginal effect of group, $F(2, 27) = 2.99$, $MSE = 0.34$, $p = .07$, $\eta_p^2 = .18$ (see Figure 3). Tukey HSD post-hocs showed that the TransRot group were significantly more accurate than the Visual-only group ($p = .02$), but the other pairwise comparisons were not significant. Transformed back to degrees, the mean errors were 17° (TransRot), 25° (Rotate) and 33° (Visual-only).

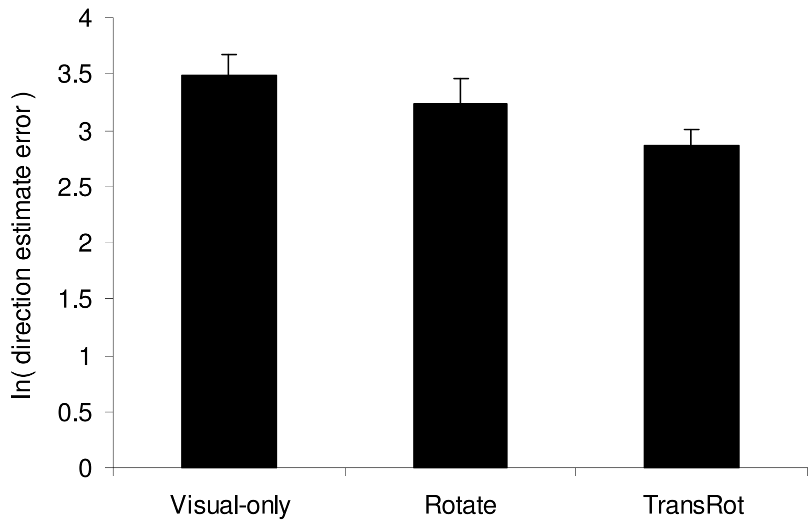

Fig. 3. Natural logarithm of participants' mean direction estimate error in Experiment 1. Error bars show standard error of the mean.

Participants' estimates of straight line distance were correlated with the actual distances, transformed to Fisher's z' to normalize the data, and analyzed using a univariate ANOVA. There was a main effect of group, $F(2, 27) = 3.31$, $MSE = 0.13$, $p = .05$, $\eta_p^2 = .20$ (see Figure 4). Tukey HSD post-hocs indicated that the TransRot group made significantly more accurate estimates than the Rotate group ($p = .02$), but the other pairwise comparisons were not significant. Transformed back to Pearson's r, the overall mean correlations were .78 (TransRot), .57 (Rotate) and .62 (Visual-only). This shows that the components of body-based information that were provided had a significant effect on participants' knowledge of the relative distances between the targets.

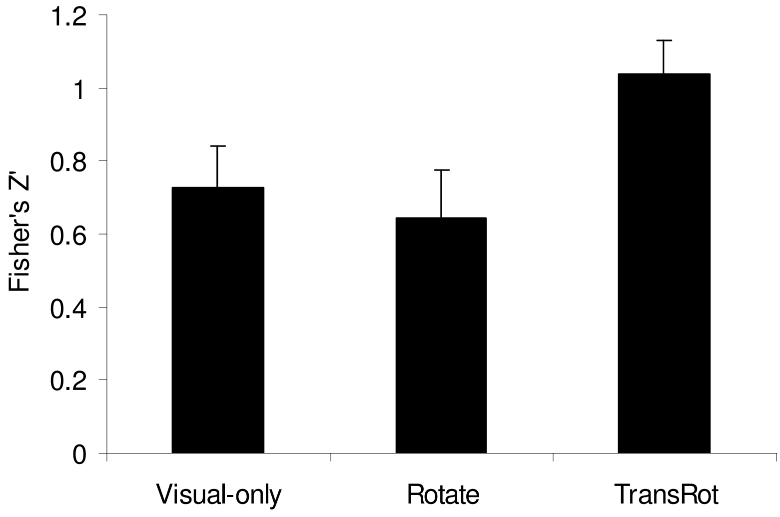

Fig. 4. Participants' mean Fisher's Z′ for estimates of relative straight line distance in Experiment 1. Error bars show standard error of the mean.

In all three groups, participants could vary their speed of travel. In the Rotate and Visual-only groups the maximum possible speed was 0.9 m/s (maximum deflection of the joystick), whereas the TransRot group's speed was not limited. Each participant's speed was averaged over 1 second intervals (this suppressed the effect of any sudden head movements) and, for all intervals with a speed > 0.25 m/s (periods when participants were not almost stationary), an average "moving" speed for the trial was calculated. Overall average moving speeds were 0.5 m/s (TransRot group) and 0.6 m/s (Rotate and Visual-only groups).

3.3 Discussion

No previous research has used such a complex task to investigate the effect of translational vs. rotational body-based information on navigation in a large-scale environment. There was a common pattern of results across all the metrics that were used, although it should be noted that the main effect for participants' direction estimates was only marginally significant due to a lack of statistical power.

When both components of body-based information were provided (the TransRot group), participants explored the marketplace more efficiently and developed a more accurate cognitive map. In Trial 1, when participants had no prior knowledge of the marketplace's layout, the Rotate and Visual-only groups traveled more than 25% further than the TransRot group, indicating that even during initial navigation of the marketplace the TransRot group developed better knowledge of where they had been, so they could concentrate on checking parts of the marketplace that they had not previously visited (Note: the fact that the targets lay in all four quadrants of the environment reduced the role that chance played in participants' explorations). These results are consistent with the advantage that full body-based information provided when participants searched a small-scale space (a 3 × 3 m room) for targets [Ruddle and Lessels 2009]. In Trial 2 the Rotate/Visual-only groups traveled more than 50% further than the TransRot group, showing that the TransRot group's knowledge of the environment continued to develop faster than that of participants in the other two groups.

The similarities between the performance of the Rotate and Visual-only groups in the present experiment are important for two reasons. First, those similarities replicate the null hypothesis (i.e., statistically insignificant) findings of previous research that has investigated navigation in large-scale/extent virtual buildings and mazes [Ruddle, Payne and Jones 1999; Ruddle and Péruch 2004]. Second, the improved performance of the TransRot group cannot have been caused by the use of a stereo display, or any general increase in "presence" that may have resulted from participants being more immersed when viewing the marketplace in an HMD, because the Rotate group used the same display as the TransRot group.

Experiment 1 had two main limitations. First, the marketplace was large-scale but only small in extent. Second, the experiment did not show whether navigational benefits occur only when both components of body-based information are provided, or whether just translational body-based information is required. Both of these limitations were addressed in Experiment 2

4. EXPERIMENT 2 (LARGE SCALE & EXTENT)

A between-participants design was used, with four groups (*TransRot*, *Translate*, *Rotate* and *Visual-only*; see Table I). The TransRot navigated by walking on the Cyberwalk omni-directional treadmill [De Luca et al. 2009], and the Translate group navigated by walking on a linear treadmill [Souman et al. 2010]. The Rotate and Visual-only groups used the same interfaces as Experiment 1.

4.1 Method

4.1.1 Participants

Forty-four individuals (21 women) with a mean age of 26 years (*SD* = 5.1) took part. All gave informed consent, took approximately 2 hours to complete the experiment, and were paid an honorarium for their participation. The study was approved by the local ethics committee.

Four participants (1 woman) withdrew because of motion sickness, two from the Translate group and one each from the Rotate and TransRot groups. The remaining participants were randomly assigned to each group, subject to the groups being gender balanced (5 men and 5 women in each).

4.1.2 Materials

To allow participants to practice the movement interface for their group, a new virtual marketplace was constructed. This measured 45 × 25 m, had 5 m wide corridors, and contained a 270 m long route that was marked with arrows and zigzagged back and forth though the marketplace. All the stalls in this marketplace were identical (there were no pictures) because its sole purpose was to let participants practice maneuvering. This was necessary for the treadmill groups, because participants needed to become familiar with the dynamics of the treadmill control algorithms. For consistency, the maneuvering practice was also performed by the other groups.

To practice the task, participants used a 4 × 2 virtual marketplace that measured 45 × 25 m, had 5 m wide corridors and 5 × 5 m stalls, and the same structure and pictures as the 4 × 2 marketplace used in Experiment 1. The two versions of the 6 × 4 test marketplaces measured 65 × 45 m, had 5 m wide corridors and 5 × 5 m stalls (see Figure 5), and the same structure and pictures as Experiment 1's test marketplaces. In other words, the task practice and test marketplaces were scaled by a factor of 6.67 in length and width, compared with the marketplaces used in Experiment 1.

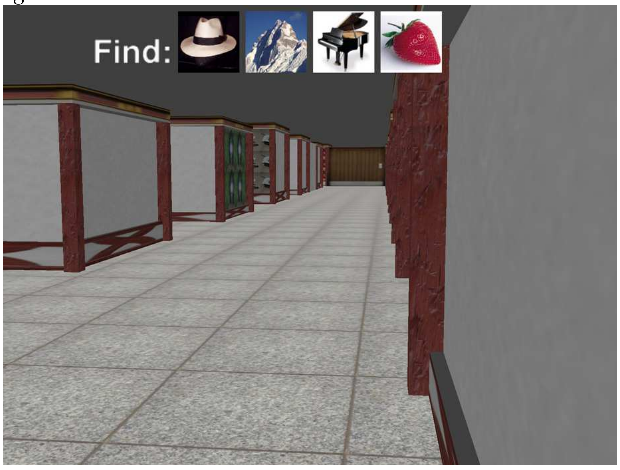

Fig. 5. Interior view of one of the test marketplaces used in Experiment 2. The layout was the same as shown in Figure 1a, but the marketplace's dimensions were 65 × 45 m.

Participants in the Rotate and Visual-only groups used the same interface as in Experiment 1, except that the joystick allowed participants to travel at up to 1.34 m/s (a faster walk than in Experiment 1), which was similar to the maximum speed of the treadmills.

The Translate participants walked on a 6 m long linear treadmill (see Figure 6a), which moved at participants' speed. Participants were tracked by a Vicon MX13 motion capture system, which provided data for the treadmill control algorithm. When participants started to walk this algorithm accelerated the treadmill belt, and decelerated it when participants slowed down or stopped (for details, see [Souman, Giordano, Frissen, De Luca and Ernst 2010]). Guide ropes that ran the length of the treadmill helped participants to stay in its centre. For safety, participants wore a harness that was attached to an overhead cable, which also supported the weight of the HMD's external video control unit. Orientation tracking was turned off, which meant that the scene displayed in the HMD did not change if participants turned their head. To look around or turn within the marketplaces, participants used the same device as the Visual-only group, which was the right joystick on a Logitech Rumblepad. This arrangement for looking around/turning was chosen because it meant that the Translate group was not provided with any rotational body-based information for their movement through the marketplaces. It is also worth noting that, in preliminary research we performed, enabling head tracking so that participants could physically look around while walking on the treadmill caused participants to travel diagonally on the treadmill surface or to make movements that triggered nausea-inducing sensory conflicts.

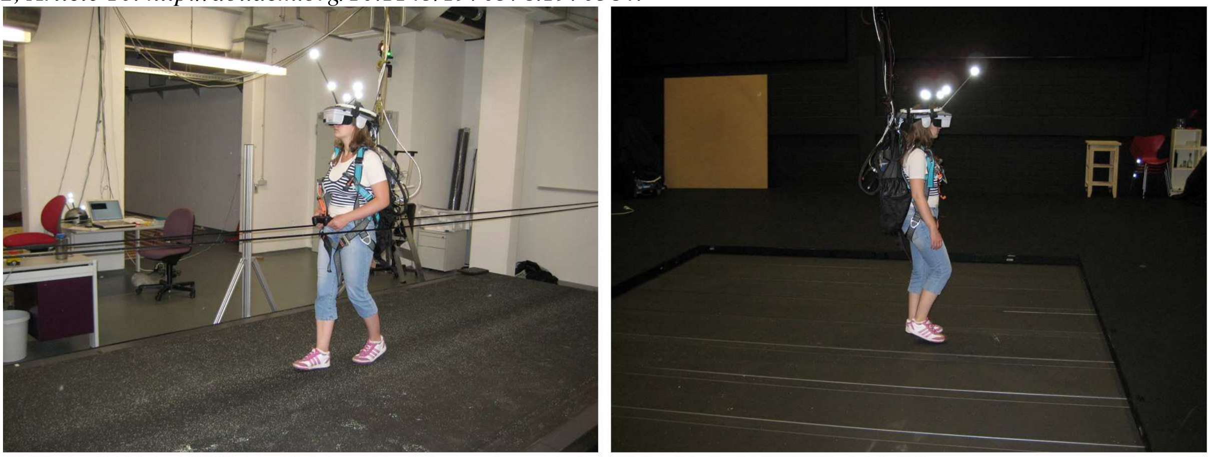

(a) (b)

Fig. 6. (a) The linear treadmill (Translate group), and (b) the Cyberwalk omni-directional treadmill (TransRot group). For illustrative purposes, the same person is shown in both photographs. The experiment used a between-participants design.

The TransRot participants walked on a 4 × 4 m omni-directional treadmill (see Figure 6b), which moved at participants' speed and used the same type of control algorithm as the linear treadmill (for details, see [De Luca, Mattone, Giordano and Bülthoff 2009]). Participants were encouraged to walk normally (the control algorithm always moved participants back toward the centre of the treadmill) and, for safety, wore a harness that was attached to an overhead cable. The cable also supported the weight of the HMD's video control unit. The engineering design of the treadmill meant that its dynamics changed slightly according to the direction in which participants traveled. To help prevent this from providing an orientation cue, the VE software oriented each marketplace one way relative to the treadmill for Trial 1, and then rotated the marketplace relative to the treadmill by 90° for Trial 2.

Participants in the TransRot, Translate and Rotate groups all wore earplugs to mask, but not totally exclude, external sounds. Safety considerations with the treadmills meant that participants needed to be able to hear instructions that the experimenter might shout in an emergency (no such episodes occurred).

4.1.3 Procedure

First, a participant practiced the task by performing two trials in the 4 × 2 marketplace. This was always done using the Visual-only configuration, so that the experimenter could explain the task face-to-face.

Next, participants practiced the movement interface for their group. The TransRot 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 practiced walking through a VE on the treadmill by making two traversals of the defined 270 m route (see above). The Translate 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 practiced walking through a VE on the treadmill by making two traversals of the 270 m route. The Rotate and Visual-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 VE by making two traversals of the 270 m route using the interface for their respective groups.

After this, participants performed two test trials in one version of the 6 × 4 test marketplace, using the interface for their group. The procedure was the same as in Experiment 1. In Trial 1, a participant searched for the targets in any order and then returned to the start point. In Trial 2, the participant again searched for the targets in any order and, at each target, estimated the direction and straight line distance to every other target and the start point. Direction estimates were performed by either physically turning (TransRot and Rotate groups) or using the joystick to turn (Translate and Visual-only groups) until the participant judged that they were looking through the stalls directly toward the specified target's location in the VE. Distance estimates were reported verbally. Once both trials were complete, the participant answered a short questionnaire.

4.2 Results

The data were analyzed using ANOVAs that had two between-participants factors (translational × rotational body-based information) and, for the distance traveled, one within-participants factor (trial). There were no significant interactions.

The distribution of the distance traveled data was normalized using a natural logarithm transformation. An ANOVA showed that participants traveled significantly less distance in Trial 2 than Trial 1, $F(1, 36) = 13.94$, $MSE = 0.06$, $p = .001$, $\eta_p^2 = .28$ (see Figure 7). However, there was no main effect for translational body-based information, $F(1, 36) = 2.15$, $MSE = 0.14$, $p = .15$, $\eta_p^2 = .06$, or rotational body-based information, $F(1, 36) = 0.04$, $MSE = 0.14$, $p = .83$, $\eta_p^2 < .01$. Six Translate, four Rotate, two TransRot and one Visual-only participant visited the targets in an identical order in both trials, and one Translate and one Visual-only participant reversed the order. Transformed back from logarithms and expressed as multiples of the shortest path, the distances for Trials 1 and 2 were 1.5 and 1.2 (TransRot), 1.6 and 1.2 (Translate), 1.7 and 1.5 (Rotate), and 1.6 and 1.4 (Visual-only).

The questionnaire showed that 10 participants played computer games frequently (at least once a week). In terms of total distance traveled they were ranked as follows within their group (overall): TransRot group 1st (2nd), 3rd (9th) and 7th (23rd), Translate group 1st (3rd), 2nd (4th) and 8th (30th), Rotate group 2nd (10th) and 3rd (12th), and Visual-only group 3rd (8th), 6th (22nd).

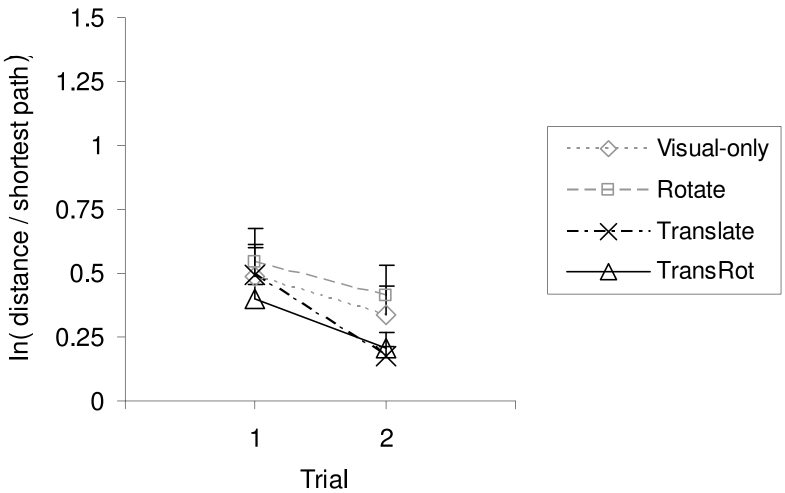

Fig. 7. Natural logarithm of the distance participants traveled during Experiment 2, expressed as multiples of the shortest path (for the shortest path, ln(distance) = 0.0). Error bars show standard error of the mean.

From participants' direction estimates, the mean absolute angular error was calculated and then the distribution of these data was normalized using a natural logarithm transformation. An ANOVA showed that participants who were provided with translational body-based information (the Translate & TransRot groups) made significantly more accurate direction estimates, $F(1, 36) = 7.54$, $MSE = 0.31$, $p = .009$, $\eta_p^2 = .17$, but there was no main effect for rotational body-based information, $F(1, 36) = 0.05$, $MSE = 0.31$, $p = .82$, $\eta_p^2 < .01$ (see Figure 8). Transformed back to degrees, the mean errors were 17° (TransRot), 19° (Translate), 32° (Rotate) and 26° (Visual-only).

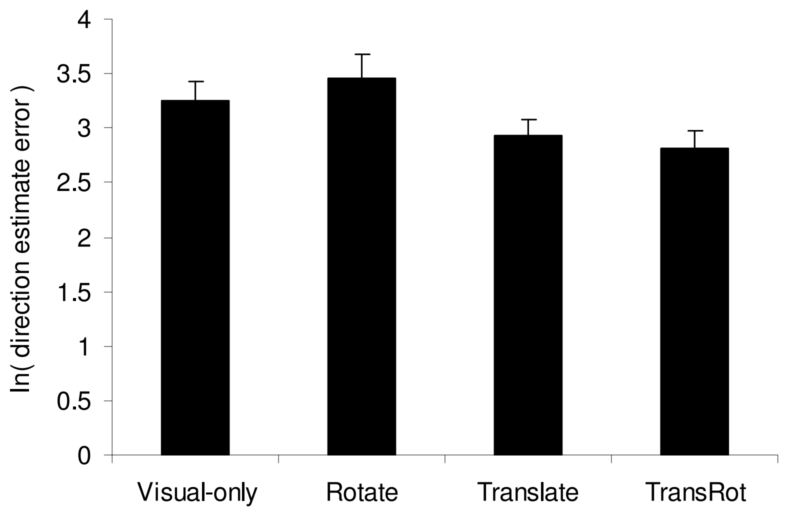

Fig. 8. Natural logarithm of participants' mean direction estimate error in Experiment 2. Error bars show standard error of the mean.

Participants' estimates of straight line distance were correlated with the actual distances, and transformed to Fisher's z' to normalize the data. An ANOVA showed that participants who were provided with translational body-based information (the Translate & TransRot groups) made significantly more accurate estimates, $F(1, 36) = 6.68$, $MSE = 0.13$, $p = .01$, $\eta_p^2 = .16$, but there was no main effect for rotational body-based information, $F(1, 36) = 0.06$, $MSE = 0.13$, $p = .80$, $\eta_p^2 < .01$ (see Figure 9). In other words, translational body-based information improved participants' knowledge of the relative distances between targets, but rotational body-based information did not. Transformed back to Pearson's r, the overall mean correlations were .74 (TransRot), .60 (Translate), .41 (Rotate) and .56 (Visual-only).

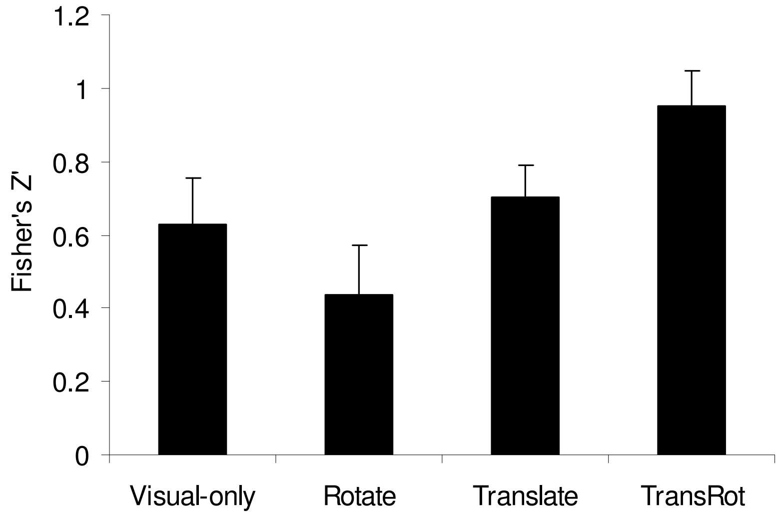

Fig. 9. Participants' mean Fisher's Z′ for estimates of relative straight line distance in Experiment 2. Error bars show standard error of the mean.

Participants' average speed while moving was calculated in the same way as Experiment 1. Overall averages were 1.2 m/s for all groups except the Translate group (0.7 m/s). That group's low average is likely to have been caused by the awkwardness of the interface (see Discussion).

4.3 Discussion

Experiment 2 was designed to answer two questions: (a) Does translational body-based information on its own provide a navigational benefit, or are both components of body-based information required? and (b) How is the navigational benefit affected by an environment that is large in both scale and extent?

The results were unequivocal. Only translational body-based information was required, and that significantly improved the accuracy of participants' cognitive map, but not participants' navigational performance in initial exploration (Trial 1) and revisiting (Trial 2) tasks. These findings have important theoretical and applied implications, which are discussed in the next section.

Although the linear and omni-directional treadmills that were used in this experiment are best of breed devices, it would inappropriate to claim that they make VE navigation the same as walking in the real world. Walking on the Cyberwalk omni-directional treadmill is sometimes likened to walking on a ship in rough seas, which results in a slightly staggering gait. With the linear treadmill, participants could only look around and turn with the gamepad joystick and, to minimize sensory conflicts, were advised to keep their head still. Even after training this interface was somewhat awkward to use, which may explain the lower accuracy of the Translate group's distance estimates compared with the TransRot group's (see Figure 9). That said, it should be emphasized that, for both the straight line distance and direction estimate data, translational body-based information caused a very significant improvement in the accuracy of participants' estimates, but rotational body-based information had no effect and there were no significant interactions.

5. GENERAL DISCUSSION

This is, by far, the most revealing study there has been into the effect that full vs. reduced body-based information has on navigation in a large-scale space: Nobody has previously attempted to study the effect the two components of body-based information have on navigation in environments that have such a high degree of ecological validity in terms of scale, extent and richness of the visual scene. The key finding was that the addition of translational body-based information (the TransRot & Translate groups) significantly improved participants' cognitive maps, whereas rotational body-based information provided no benefit. A notable secondary finding was that, as spatial extent decreased, body-based information provided a large additional benefit for navigational performance (the distance that participants traveled). The study's findings are summarized in Table II.

Table II. Summary of the environments, main effects of body-based information (experiment group or translational/rotational component), and significant post-hoc tests in the experiments (NS indicates $p > .05$; * indicates $p \leq .05$; ** indicates $p \leq .01$). The effect sizes of the significant effects were all small ($0.1 < \eta_p^2 < 0.3$).

Environment/metric		**Experiment 1**	**Experiment 2**
Environment	Scale	Large	Large
	Extent	Small	Large
Metric	Distance traveled	Group** TransRot vs. Rotate post-hoc** TransRot vs. Visual-only post-hoc*	Translational component NS Rotational component NS
	Direction estimates	Group NS TransRot vs. Visual-only post-hoc*	Translational component** Rotational component NS
	Distance estimates	Group* TransRot vs. Rotate post-hoc*	Translational component** Rotational component NS

5.1 Theoretical implications

Theories of spatial knowledge acquisition tend to discount the contribution that body-based information makes when people perform complex navigational tasks in everyday settings. The reasons are twofold. First, everyday settings contain a rich assortment of visual information, and that alone is sufficient for people to accurately judge the angles they turn through and relative distances they travel [Bremmer and Lappe 1999; Riecke, van Veen and Bülthoff 2002; Sun et al. 2004], so it is assumed that body-based information isn't needed. Second, path integration errors accumulate over time [Loomis, Klatzky, Golledge and Philbeck 1999; Souman et al. 2009] and so body-based information is assumed to become less and less important as spatial extent increases. However, the results of the present study show that these assumptions are flawed, because translational body-based information clearly improved the accuracy of participants' cognitive maps, even when the environment had a grid-like structure which would have allowed participants to quantify distance with a visual counting strategy (the number of blocks traversed). It should also be noted that when visual scenes lack a rich assortment of landmarks then the brain gives body-based information a greater weight than visual information for distance estimation [Campos et al. 2010] and the provision of body-based information reduces navigational variance in triangle completion tasks [Kearns et al. 2002].

Related to this, we propose two hypotheses that need to be investigated in future research. The first is that translational body-based information significantly improves people's ability to take shortcuts, because successful shortcutting involves traveling an appropriate distance in the correct direction. Doing so requires an accurate cognitive map, and particularly knowledge of directions and relative distances that were used as cognitive map metrics in the present study. Second, translational body-based information is likely to have an even greater benefit in environments such as towns that have irregular layouts, because counting strategies are more difficult to use.

Studies conducted using rats, some of which have even required the rats to navigate through VEs on an omni-directional treadmill [Hölscher et al. 2005], have revealed the neuronal mechanisms that are involved in large scale navigation (place cells, grid cells, etc.), and shown that body-based information increases neuronal (theta wave) activity (for reviews, see [Brotons-Mas et al. 2006; McNaughton et al. 2006]). These studies rely on invasive techniques that are not possible with humans, but imaging technology is now suggesting that the same neuronal mechanisms are used in human navigation [Doeller et al. 2010]. Challenges for the future are to: (a) understand how body-based information increases neuronal activity in humans, and (b) develop a theoretical model that shows how changes in this activity lead to the type of improvements in the accuracy of human's cognitive maps that were found in the present study.

Theories of spatial knowledge also need to take greater account of the extent of an environment. At present, extent is only used to distinguish between spaces within, around (arms' length) or beyond a person's body [Tversky et al. 1999], with the latter classified as small- or large-scale, depending on whether people need to travel through a space to see its layout. However, the present study shows that extent is also an important attribute of large-scale spaces. The overall improvement in participants' navigational performance from Experiment 1 to Experiment 2 can be attributed to participants taking more care in planning where to travel, because the time cost of an error increased with spatial extent. Similar results occurred when participants navigated VEs [Ruddle, Howes, Payne and Jones 2000] and graphical menus [O'Hara and Payne 1998]. The effect of body-based information on participants' navigational performance, which was only significant in Experiment 1, is likely to have been caused by either the cognitive cost of maneuvering in narrow corridors or the smaller amount of time for which each visual cue was visible. Referring to the latter, if a participant traveled down the centre of a corridor at the Visual-only and Rotate groups' maximum speed then the participant would have traveled past a stall in 1.7 seconds in Experiment 1, compared with 7.5 seconds in Experiment 2. The small amount of time that the picture in stalls were visible in Experiment 1 is likely to have increased the role that body-based information played in remembering where one had traveled. The same is likely to be true for time-pressured applications that involve larger environments, for example using VEs to train emergency evacuation procedures.

Roy A. Ruddle, Ekaterina Volkova, & Heinrich H. Bülthoff. (2011). Walking improves your cognitive map in environments that are large-scale and large in extent. ACM Transactions on Computer-Human Interaction, 18, 2, Article 10. http://doi.acm.org/10.1145/1970378.1970384.

5.2 Applications

Humans are very good at learning new environments in the real world because we develop a cognitive map, as well as route knowledge, from the outset [Montello 1998]. However, in VEs a substantial minority of participants experienced great difficulty learning moderately complex layouts [Ruddle 2001]. For many years, researchers have sought to solve these navigational problems by developing interfaces that would allow users to physically walk through large VEs. Now, for the first time in environments that are large in extent as well as large-scale, there is evidence that walking interfaces do significantly improve certain aspects of participants' spatial knowledge.

The three types of physical walking interface that are suitable for large VEs are: (a) treadmills [Darken et al. 1997; De Luca, Mattone, Giordano and Bülthoff 2009; Hollerbach et al. 2003], (b) walking-in-place algorithms [Feasel et al. 2008; Slater et al. 1995; Templeman et al. 1999], and (c) redirected walking techniques (for a review, see [Peck et al. 2009]). How do the present study's findings guide the implementation of these interfaces?

Omni-directional treadmills were developed to allow people to walk "normally" through VEs of unlimited size, but are large, heavy and extremely specialized pieces of equipment (the Cyberwalk treadmill weighs 12 tonnes). However, the present study's key finding is that only translational body-based information is required, so the physical turning capability of an omni-directional treadmill is not needed. In other words, a linear treadmill is sufficient.

Although a linear treadmill provides sufficient body-based information, it is awkward to use with an HMD because all turning has to be performed using an abstract device such as a joystick. However, a straightforward solution would be to combine the treadmill with a wide FOV projected display (e.g., a curved theatre display [Trutoiu et al. 2009]), so participants could glance around using normal head/eye movements, turn using a joystick, but be able to walk straight because the treadmill itself is visible. This also removes the general encumbrance of wearing an HMD.

An alternative that deserves investigation is to combine walking-in-place with a wide FOV projected or CAVE display. The results of the present study suggest that, for general travel through a space, the walking-in-place algorithm would only need to detect "forward" motion, although this would not satisfy the requirements of applications that necessitate rapid maneuvering movements, such as those that occur in some military situations [Whitton et al. 2005].

Redirected walking techniques depend on reorienting users without them realizing, so they can walk forever despite being in a limited physical space. The fact that rotational body-based information had no effect on participants' navigation in the present study, or any of our previous research, suggests that the reorienting process in redirected walking is not likely to inhibit users' performance. However, current versions of redirected walking depend on artificially diverting users' attention during reorientation [Peck, Fuchs and Whitton 2009], which would be unrealistic for most applications. Overall, therefore, the linear treadmill or walking-in-place solutions are likely to be more appropriate.

Finally, spatial metaphors have often been proposed for the organization of information in a variety of applications. For example, spatial layouts are part of every mainstream desktop interface today (Windows, Mac

& Linux), and spatial hypertext systems have been developed for use in areas such as digital libraries [Buchanan et al. 2004]. Some user studies have shown that a spatial interface is beneficial [Robertson et al. 1998], but more recent studies that used the same general task and layouts have not [Cockburn and McKenzie 2002]. In those studies, documents were laid out as if on a desk, which in spatial cognition terms is a small-scale space. The present study highlighted fundamental differences between the sensory information that is useful for navigating large-scale spaces, compared with earlier studies that used small-scale spaces. Therefore, it may be more beneficial to use the metaphor of a large-scale space to organize information, for example, portraying the content and structure of a website as streets and buildings on a city map [Ruddle 2010].

PRIOR PUBLICATION

This article describes research that is entirely new and has never been published before.

ACKNOWLEDGEMENTS

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- 2008-000-10008-0). We thank Michael Kerger, Joachim Tesch and Betty Mohler for their technical assistance.

REFERENCES

AVRAAMIDES, M.N., KLATZKY, R.L., LOOMIS, J.M. AND GOLLEDGE, R.G. 2004. Use of cognitive versus perceptual heading during imagined locomotion depends on the response mode. *Psychological Science 15*, 403-408.

BREMMER, F. AND LAPPE, M. 1999. The use of optical velocities for distance discrimination and reproduction during visually simulated self motion. *Experimental Brain Research 127*, 33-42.

BROOKS, F.P., AIREY, J., ALSPAUGH, J., BELL, A., BROWN, R., HILL, C., NIMSCHECK, U., RHEINGANS, P., ROHLF, J., SMITH, D., TURNER, D., VARSHNEY, A., WANG, Y., WEBER, H. AND YUAN, X. 1992. Six generations of building walkthrough: Final technical report to the National Science Foundation. In *Technical Report #TR92-026* Department of Computer Science, University of North Carolina, Chapel Hill, NC.

BROTONS-MAS, J.R., O'MARA, S. AND SANCHEZ-VIVES, M.V. 2006. Neural processing of spatial information: What we know about place cells and what they can tell us about presence. *Presence: Teleoperators and Virtual Environments 15*, 485-499.

BUCHANAN, G., BLANDFORD, A., THIMBLEBY, H. AND JONES, M. 2004. Integrating information seeking and structuring: Exploring the role of spatial hypertexts in a digital library. In *Proceedings of the 15th ACM Conference on Hypertext and Hypermedia* ACM, New York, 225-234.

CAMPOS, J.L., BYME, P. AND SUN, H.J. 2010. The brain weights body-based cues higher than vision when estimating walked distances. *European Journal of Neuroscience 31*, 1889-1898.

CHANCE, S.S., GAUNET, F., BEALL, A.C. AND LOOMIS, J.M. 1998. Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments 7*, 168-178.

COCKBURN, A. AND MCKENZIE, B. 2002. Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* ACM, New York, 203-210.

DARKEN, R., COCKAYNE, W. AND CARMEIN, D. 1997. The omni-directional treadmill: A locomotion device for virtual worlds. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology* ACM, New York, 213-221.

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.

DOELLER, C.F., BARRY, C. AND BURGESS, N. 2010. Evidence for grid cells in a human memory network. *Nature 463*, 657-661.

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.

FOO, P., WARREN, W.H., DUCHON, A. AND TARR, M.J. 2005. Do humans integrate routes into a cognitive map? Map- versus landmark-based navigation of novel shortcuts. *Journal of Experimental Psychology: Learning, Memory and Cognition 31*, 195-215.

GRANT, S.C. AND MAGEE, L.E. 1998. Contributions of proprioception to navigation in virtual environments. *Human Factors 40*, 489-497.

HOLLERBACH, J., CHECCACCI, D., NOMA, H., YANAGIDA, Y. AND TETSUTANI, N. 2003. Simulating side slopes on locomotion interfaces using torso forces. In *Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environments and Teleoperation* IEEE, Los Alamitos, CA, 91-98.
HÖLSCHER, C., SCHNEE, A., DAHMEN, H., SETIA, L. AND MALLOT, H.A. 2005. Rats are able to navigate in virtual environments. *Journal of Experimental Biology 208*, 561-569.
JANSEN-OSMANN, P. AND FUCHS, P. 2006. Wayfinding behavior and spatial knowledge of adults and children in a virtual environment: The role of landmarks. *Experimental Psychology 5*, 171-181.
KEARNS, M.J., WARREN, W.H., DUCHON, A.P. AND TARR, M.J. 2002. Path integration from optic flow and body senses in a homing task. *Perception 31*, 349-374.
KLATZKY, R.L., LOOMIS, J.M., BEALL, A.C., CHANCE, S.S. AND GOLLEDGE, R.G. 1998. Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science 9*, 293-298.
LESSELS, S. AND RUDDLE, R.A. 2004. Changes in navigational behaviour produced by a wide field of view and a high fidelity visual scene. In *Proceedings of the 10th Eurographics Symposium on Virtual Environments* Eurographics Association, Aire-la-Ville, Switzerland, 71-78.
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.
MCNAUGHTON, B.L., BATTAGLIA, F.P., JENSEN, O., MOSER, E.I. AND MOSER, M.-B. 2006. Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience 7*, 663-678.
MONTELLO, D.R. 1991. The measurement of cognitive distance: Methods and construct validity. *Journal of Environmental Psychology 11*, 101-122.
MONTELLO, D.R. 1998. A new framework for understanding the acquisition of spatial knowledge in large-scale environments. In *Spatial and temporal reasoning in geographic information systems*, R. GOLLEDGE AND M. EGENHOFER Eds. Oxford University, New York, 143-154.
O'HARA, K. AND PAYNE, S. 1998. The effects of operator implementation cost on planfulness of problem solving and learning. *Cognitive Psychology 35*, 34-70.
O'KEEFE, J. AND NADEL, L. 1978. *The hippocampus as a cognitive map*. Oxford University, Oxford.
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.
RIECKE, B.E., VAN VEEN, H.A.H.C. AND BÜLTHOFF, H.H. 2002. Visual homing is possible without landmarks: A path integration study in virtual reality. *Presence: Teleoperators and Virtual Environments 11*, 443-473.
ROBERTSON, G., CZERWINSKI, M., LARSON, K., ROBBINS, D.C., THIEL, D. AND VAN DANTZICH, M. 1998. Data mountain: Using spatial memory for document management. In *Proceedings of the 11th annual ACM Symposium on User Interface Software and Technology* ACM, New York, 153-162.
RUDDLE, R.A. 2001. Navigation: Am I really lost or virtually there? In *Engineering psychology and cognitive ergonomics*, D. HARRIS Ed. Ashgate, Burlington, VT, 135-142.
RUDDLE, R.A. 2010. INSPIRE: A new method for mapping information spaces. In *Proceedings of the 14th International Conference on Information Visualisation (IV 2010)* IEEE, Los Alamitos, CA, 273-279.
RUDDLE, R.A., HOWES, A., PAYNE, S.J. AND JONES, D.M. 2000. The effects of hyperlinks on navigation in virtual environments. *International Journal of Human-Computer Studies 53*, 551-581.
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., PAYNE, S.J. AND JONES, D.M. 1997. Navigating buildings in "desk-top" virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied 3*, 143-159.
RUDDLE, R.A., PAYNE, S.J. AND JONES, D.M. 1999. Navigating large-scale virtual environments: What differences occur between helmet-mounted and desk-top displays? *Presence: Teleoperators and Virtual Environments 8*, 157-168.
RUDDLE, R.A. AND PÉRUCH, P. 2004. Effects of proprioceptive feedback and environmental characteristics on spatial learning in virtual environments. *International Journal of Human-Computer Studies 60*, 299-326.
RUDDLE, R.A., VOLKOVA, E., MOHLER, B. AND BÜLTHOFF, H.H. in press. The effect of landmark and body-based sensory information on route learning. *Memory and Cognition*.
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.
SOUMAN, J.L., FRISSEN, I., SREENIVASA, M.N. AND ERNST, M.O. 2009. Walking straight into circles. *Current Biology 19*, 1538-1542.
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.
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.
SUN, H.J., CAMPOS, J.L., YOUNG, M., CHAN, G.S.W. AND ELLARD, C.G. 2004. The contributions of static visual cues, nonvisual cues, and topic flow in distance estimation. *Perception 33*, 49-65.
TAN, D.S., GERGLE, D., SCUPELLI, P. AND PAUSCH, R. 2006. Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction 13*, 71-99.
TEMPLEMAN, J., DENBROOK, P. AND SIBERT, L. 1999. Virtual locomotion: Walking in place through virtual environments. *Presence: Teleoperators and Virtual Environments 8*, 598-617.
THORNDYKE, P. AND HAYES-ROTH, B. 1982. Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology 14*, 560-589.
TRUTOIU, L.C., MOHLER, B.J., SCHULTE-PELKUM, J. AND BÜLTHOFF, H.H. 2009. Circular, linear, and curvilinear vection in a large-screen virtual environment with floor projection *Computer & Graphics 33*, 47-58.

TVERSKY, B., MORRISON, J.B., FRANKLIN, N. AND BRYANT, D.J. 1999. Three spaces of spatial cognition. *Professional Geographer 51*, 516-524.
WALLER, D., HUNT, E. AND KNAPP, D. 1998. The transfer of spatial knowledge in virtual environment training. *Presence: Teleoperators and Virtual Environments 7*, 129-143.
WALLER, D., LOOMIS, J.M. AND HAUN, D.B.M. 2004. Body-based senses enhance knowledge of directions in large-scale environments. *Psychonomic Bulletin & Review 11*, 157-163.
WEATHERFORD, D.L. 1985. Representing and manipulating spatial information from different environments: Models to neighborhoods. In *The development of spatial cognition* R. COHEN Ed. Erlbaum, Hillsdale, New Jersey, 41-70.
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.
WITMER, B., BAILEY, J., KNERR, B. AND PARSONS, K. 1996. Virtual spaces and real-world places: Transfer of route knowledge. *International Journal of Human-Computer Studies 45*, 413-428.
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.