

Spatial cognitive implications of teleporting through virtual environments

Lucia A. Cherep, Alex F. Lim, Jonathan W. Kelly, Devi Acharya, Alfredo Velasco,

Emanuel Bustamante, Alec G. Ostrander, Stephen B. Gilbert

Iowa State University

Author Note

Pre-registrations for Experiments 3-5, as well as videos, data, supplementary analyses, and experiment code are available on the Open Science Framework (<https://osf.io/m4zfv/>). This work was supported by a grant from the National Science Foundation (CHS-1816029 to J.W.K. and S.B.G.).

Correspondence concerning this article should be addressed to Jonathan Kelly, Department of Psychology, Iowa State University, Ames, IA, 50011-1041. Contact: jonkelly@iastate.edu.

Abstract

Teleporting is a popular interface to allow virtual reality users to explore environments that are larger than the available walking space. When teleporting, the user positions a marker in the virtual environment and is instantly transported without any self-motion cues. Five experiments were designed to evaluate the spatial cognitive consequences of teleporting, and to identify environmental cues that could mitigate those costs. Participants performed a triangle completion task by traversing two outbound path legs before pointing to the unmarked path origin.

Locomotion was accomplished via walking or two common implementations of the teleporting interface distinguished by the concordance between movement of the body and movement through the virtual environment. In the partially concordant teleporting interface, participants teleported to translate (change position) but turned the body to rotate. In the discordant teleporting interface, participants teleported to translate and rotate. Across all 5 experiments, discordant teleporting produced larger errors than partially concordant teleporting which produced larger errors than walking, reflecting the importance of translational and rotational self-motion cues. Furthermore, geometric boundaries (room walls or a fence) were necessary to mitigate the spatial cognitive costs associated with teleporting, and landmarks were helpful only in the context of a geometric boundary.

Keywords: Spatial cognition; Navigation; Triangle completion; Teleporting; Virtual reality

Spatial cognitive implications of teleporting through virtual environments

Virtual reality (VR) is becoming increasingly popular in industry, education, and entertainment, partly due to the availability of low-cost systems such as the Oculus Rift and HTC Vive (Mainelli, Shirer, & Ubrani, 2019). One especially compelling feature of most modern VR systems is the ability to explore virtual environments (VEs) by physically walking and turning, whereby movement of the user is tracked by cameras and accelerometers within a finite tracked space and the graphics in the head-mounted display (HMD) are updated based on the tracked movement. However, VEs commonly exceed the size of the walkable tracked space, and therefore require a locomotion interface that allows for movement beyond that space. Perhaps the most common locomotion interface to allow for exploration of large VEs is teleportation (Figure 1), whereby the user selects a location on the ground plane and is immediately transported to that location without any self-motion cues.



Figure 1. Examples of two different teleporting interfaces from virtual reality games (left: The Lab, Valve Corporation; right: Robo Recall, Epic Games).

The teleportation interface is now widespread in VR applications, most likely due to its ease of use (Bozgeyikli, Raij, Katkooori, & Dubey, 2016; Langbehn, Lubos, & Steinicke, 2018) and reduced likelihood of sickness compared with other interfaces that include visual self-motion (i.e., optic flow) (Christou & Aristidou, 2017; Moghadam, Banigan, & Ragan, 2018; Langbehn et al., 2018; Weißker, Kunert, Fröhlich, & Kulik, 2018). Yet, the popularity of teleportation interfaces may come at a spatial cognitive cost. In particular, the lack of self-motion cues when teleporting may disrupt spatial updating, the process of keeping track of self-location during travel. Spatial updating failure can lead to disorientation (i.e., not knowing self-location relative to an external reference point or reference direction), which can be corrected only by using external piloting cues (e.g., landmarks) to reorient. Accurate spatial updating is critical for a successful VE for non-gaming purposes, particularly in time-sensitive domains, such as remote medical assistance and training for extreme workplace circumstances, situations in which disorientation could seriously hamper progress toward the user's goals. The current study investigated the influence of two common teleporting interfaces on spatial updating in VEs with and without piloting cues.

Spatial Updating

Spatial updating (updating self-location during travel) depends critically on self-motion cues, which can be broadly categorized as internal cues and external cues (see Chrastil & Warren, 2012). Internal self-motion cues include vestibular stimulation, proprioception, and efference copies of motor commands, and are herein referred to collectively as body-based self-motion cues, or body-based cues for ease of exposition. External self-motion cues are provided by optic and acoustic flow that occur during self-motion. Spatial updating is informed by both types of self-motion cues, although research points to a particularly important role for body-

based cues (Chance, Gaunet, Beall, & Loomis, 1998; Grant & Magee, 1998; Ruddle, Volkova, & Bühlhoff, 2011; Ruddle, Volkova, Mohler, & Bühlhoff, 2011; Waller, Loomis, & Haun, 2004).

Triangle completion is a prototypical spatial updating task in which the participant moves along two legs of an outbound path before pointing to or directly returning to the path origin. In a foundational study (Klatzky, Loomis, Beall, Chance, & Golledge, 1998), participants wore a head-mounted display (HMD) and performed triangle completion in a featureless VE consisting only of a grassy field devoid of landmarks. Three relevant conditions manipulated the extent to which body-based self-motion cues were available, whereas visual self-motion (i.e., optic flow) through the VE was experienced in all conditions. Triangle completion error was greatest when outbound path movement was purely visual, comparatively small when participants physically walked and turned, and equally small when participants physically rotated through the turn separating the two path legs but received only optic flow when traveling along the two straight path legs. These results indicate that successful spatial updating requires body-based cues associated with rotation (change in orientation), but body-based cues associated with translation (change in position) are unnecessary. These findings echo research on imagined spatial updating, which is more difficult when the imagined movement involves rotation compared to translation (May, 2004; Presson & Montello, 1994; Rieser, 1989).

In contrast to the findings reported by Klatzky et al. (1998), other research points to an important role for translational body-based cues. In one study (Ruddle & Lessels, 2006), participants performed a foraging task that involved exploring a small VE while peering into several boxes to find search targets. Researchers manipulated the extent to which body-based self-motion cues were available during search. The primary dependent measure was the number of times a box was checked twice, which could indicate disorientation due to failure of spatial

updating. Similar to the findings of Klatzky et al., errors were highest when locomotion was purely visual (controlled solely by joystick manipulation) and lowest when locomotion involved physical walking and turning. In contrast to the results of Klatzky et al., errors when participants physically rotated but received only visual information about translation were just as high as in the purely visual condition, indicating an important role for body-based translational cues. These results indicate that successful spatial updating requires body-based cues associated with translation and rotation and that rotational body-based cues are insufficient.

Although there is general consensus in past research that body-based cues are important to spatial updating, there is disagreement as to whether the critical contribution comes from rotational cues or translational body-based cues. It is possible that task difficulty could explain the discrepant findings, whereby complex tasks such as foraging require both translational and rotational body-based cues but simple tasks such as triangle completion only require rotational body-based cues (Ruddle, Volkova, Mohler, & Bühlhoff, 2011). The current study employed a triangle completion task, and the prediction followed that body-based cues associated with rotation would be more relevant to task performance than body-based cues associated with translation.

The role of optic flow in spatial updating appears to be smaller than that of body-based cues. Although triangle completion can sometimes be performed with optic flow alone (Riecke, van Veen, & Bühlhoff, 2002; but see Klatzky et al., 1998), spatial updating is worse with visual compared to body rotation (Wraga, Creem, & Proffitt, 2004) and the influence of optic flow is minimal when body-based cues are also present (Kearns, Warren, Duchon & Tarr, 2002).

Piloting

Separate from the process of spatial updating, locations are commonly stored in memory with respect to environmental cues such as landmarks, and guidance that is based on those landmarks is referred to as piloting. For example, one could remember a location within a landmark-rich environment by encoding the bearings and distances to landmarks while standing at that location, and then subsequently return to that location by moving to the location that produces those same landmark bearings and distances (Waller, Loomis, Golledge, & Beall, 2000). Proximal landmarks primarily provide information about position, whereas distal landmarks primarily provide information about direction (Padilla, Creem-Regehr, Stefanucci, & Cashdan, 2017). Environmental shape, such as that defined by room walls, also provides a salient cue that can be used for piloting (Kelly, McNamara, Bodenheimer, Carr, & Rieser, 2008), specifying both position and orientation of the navigator within the environment.

Cue Combination

Self-motion cues and piloting cues can be combined to improve spatial updating performance (Chen, McNamara, Kelly, & Wolbers, 2017; Sjolund, Kelly, & McNamara, 2018; Nardini, Jones, Bedford, & Bradick, 2008). For example, adult participants in a triangle completion study (Nardini et al., 2008) walked the outbound path with access to self-motion cues and proximal landmarks before attempting to return directly to the path origin. Experimental manipulations prior to executing the return path created single cue (path integration only or landmark only) and dual cue conditions. Response variance, measured over repeated trials, was lowest when both cues were available during return path execution compared to when only one cue was available and indicated statistically optimal combination of self-motion and piloting cues. Self-motion cues and room shape have also been found to be optimally combined during

spatial updating (Sjolund et al.). These results indicate that spatial updating performance will suffer without body-based self-motion cues, even when rich landmarks are available to support piloting.

Concordance: A Characteristic of Locomotion Interfaces for Virtual Environments

The most natural way to explore a VE is by walking and turning one's body. In this case, movement through the VE is **concordant** with movement of the user's body, and thus all self-motion cues are available to enable spatial updating (Figure 2, left panel). However, limited tracking space relative to the size of the VE necessitates a locomotion interface. The focus of the current study is on two implementations of the teleporting interface. The most common implementation of the teleporting interface involves teleporting to translate and using the body to rotate. To translate, the user positions a marker on the ground plane and is then translated to that position without self-motion cues (visual or body-based), but to rotate, the user must rotate the body and head. In this case, movement through the VE is **partially concordant** with movement of the user's body (Figure 2, middle panel). That is, rotational movement through the VE is concordant with body movement, but translational movement is discordant with body movement. This interface is frequently found in VR games, such as *The Lab* by Valve or *Doom VFR* by Bethesda Softworks. Other partially concordant interfaces include 1) peddling a stationary bike (Sun, Chan, & Campos, 2004), whereby limb movements are consistent with movement through the VE but vestibular cues always indicate that the user is stationary, 2) scaled translational gain, whereby physical steps are converted into larger virtual steps (Interrante, Ries, & Anderson, 2007; Williams, Narasimham, McNamara, Carr, Rieser, & Bodenheimer, 2006), and 3) redirected walking, whereby the user is redirected by modifying the relationship between rotation of the body and rotation through the VE, especially when the discrepancy is above

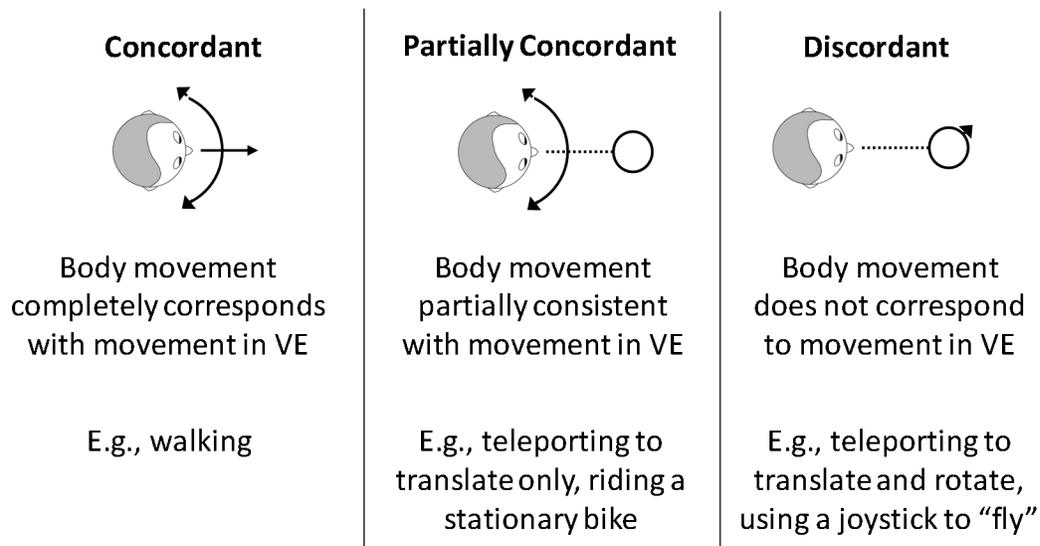


Figure 2. The concordance framework, whereby locomotion interfaces are categorized based on the extent to which movement of the user’s body corresponds with movement through the VE.

Illustrations provide examples corresponding to the three interfaces used in the current study (left panel: user physically rotates and translates to move through the VE; middle panel: user physically rotates but teleports to translate; right panel: user teleports to translate and rotate).

threshold (Grechkin, Thomas, Azmandian, Bolas, & Suma, 2016; Steinicke, Bruder, Jerald, Frenz, & Lappe, 2010). In each of those examples there is partial agreement between movement of the user and movement through the VE. Another implementation of the teleporting interface involves teleporting to translate and rotate through the VE, with no associated body movement. To rotate and translate, the user positions and orients a marker (e.g., an arrow) on the ground plane and is then teleported to that location and orientation. In this case, movement through the VE is **discordant** with movement of the user’s body (Figure 2, right panel). This interface can be seen in games such as Robo Recall by Epic Games and The Gallery Episode 1: Call of the Starseed by Cloudhead Games. Another common discordant interface is when the user manipulates a joystick or other control device to initiate smooth visual movement (i.e., to “fly”)

through the VE without accompanying body movement. However, such joystick interfaces are not favored due to their propensity to cause sickness (e.g., Christou & Aristidou, 2017). In most VEs the user can switch between multiple interfaces (e.g., walking until reaching a physical boundary and switching to teleporting), and therefore locomotion is often accomplished through a combination of several interfaces varying in concordance.

Despite the popularity of teleporting interfaces for VR, there has been relatively little research on the spatial cognitive consequences of teleporting. Existing research on the spatial cognitive effects of teleporting has typically compared teleporting with other locomotion interfaces that also lack body-based cues (Bowman, Koller, & Hodges, 1997; Christou & Aristidou, 2017; Langbehn et al., 2018; Moghadam et al., 2018; Weißker et al., 2018). Therefore, the unique contributions of the current project are to identify the spatial cognitive cost of teleporting with and without rotational body-based cues and to determine whether environmental cues, such as landmarks and spatial boundaries, can reduce potential spatial cognitive costs.

Study Overview and Predictions

Three of the locomotion interfaces described above, walking, partially concordant teleporting, and discordant teleporting, map closely onto conditions that have been tested in the spatial updating literature (e.g., Klatzky et al. and Ruddle & Lessels, 2006). This study presents the results of five experiments that manipulated access to self-motion cues as defined by the three locomotion interfaces while measuring performance on a triangle completion task. We predicted that walking would produce lower triangle completion errors compared to discordant teleporting, and that partially concordant teleporting would produce errors similar to walking since body-based cues associated with rotation have been reported to be more important to task performance than those associated with translation (Klatzky et al., 1998).

	Virtual environments	Research interest
Experiment 1	Open field Landmarks	Impact of landmarks (proximal + distal)
Experiment 2	Open field Classroom	Impact of orienting boundary (walls) with landmarks (proximal only)
Experiment 3	Open field Square fence + landmarks	Impact of orienting boundary (fence) with landmarks (proximal + distal)
Experiment 4	Square fence only Square fence + landmarks	Impact of landmarks (proximal + distal) given orienting boundary
Experiment 5	Circular fence only Circular fence + landmarks	Impact of landmarks (proximal + distal) given non-orienting boundary

Table 1. Overview of the VEs used in each experiment, and the research interest that motivated the choice of VEs. Each experiment compared triangle completion performance in two environments when using three locomotion interfaces: walking, partially concordant teleporting, and discordant teleporting.

For applications using the teleporting interface, it is important to know whether environmental cues can mitigate the potentially negative effects of teleporting on spatial updating. Therefore, we also explored the influence of piloting cues, including landmark cues and geometric (i.e., boundary shape) cues. Table 1 lists the environmental manipulations used in each experiment. We expected piloting cues to lead to lower triangle completion errors across all locomotion interfaces because piloting cues can be combined with self-motion cues when available (e.g., Nardini et al., 2008) and can be used even if participants become disoriented. We predicted that the beneficial effect of piloting cues would be greatest when using the discordant teleporting interface because of the increased chance of disorientation in the absence of all self-motion cues.

Estimates of sample size were based on two closely related studies in which rotational and translational body-based cues were manipulated (Klatzky et al., 1998; Ruddle & Lessels, 2006). Those studies used a between-participant design with 10 participants in each condition. The current study used a completely within-participant design, but also included manipulations of environmental cues that were expected to interact with manipulations of self-motion cues. Due to additional counterbalancing constraints, a sample size of 24 was used in all experiments.

Experiment 1

Experiment 1 was designed to evaluate the effects of interface concordance and piloting cues on spatial updating. Participants completed a triangle completion task in VR using three navigation interfaces: walking, partially concordant teleporting, and discordant teleporting. The VE was an endless grassy field with or without landmarks.

Method

Participants.

Twenty-six students (9 men, 17 women) at Iowa State University participated in exchange for course credit. Data from two participants were removed (see Results) leaving 24 total participants (8 men, 16 women) in the dataset.

Hardware and software.

The VEs were displayed on an HTC Vive HMD. Graphics displayed in the Vive were generated on a Windows 10 computer with an Intel 6700K processor and Nvidia GeForce GTX 1070 graphics card. Unity software displayed stereoscopic images at 1080×1200 resolution per eye with 100° horizontal \times 110° vertical binocular field of view. Images refreshed at a rate of 90 Hz and reflected head position and orientation tracked by the Lighthouse tracking system sold

with the Vive. One wireless handheld controller, also sold with the Vive, was used by participants to control the teleporting interfaces and to respond on each trial.

Stimuli.

Videos showing the triangle completion task with each interface and in each environment are available on the Open Science Framework: <https://osf.io/m4zfv/>. The open field VE (Figure 3, top-left) consisted of an infinite ground plane with grass texture and uniform blue sky. The landmarks VE (Figure 3, top-middle) included proximal and distal landmarks in addition to the grass field with blue sky. Proximal landmarks (e.g., plants, trees, and bench) were located near the intended triangle paths but were placed far enough away to not be directly on path. Distal landmarks (bridge, building, mountains, and arch) were placed in the far distance and were spaced every 90°.

The path used to perform the triangle completion task was marked by semi-transparent vertical posts, each 1 m tall and .25 m in diameter. The beginning of the path was marked by a green post, the end of the first path leg was marked by a yellow post, and the end of the second path leg was marked by a red post. Green post locations were located in a ring around the center of the environment. The yellow post led the participant toward the center of the environment, and the red post led the participant away from the center of the environment. At the base of each post was a white arrow with a height of 7.5 cm (shown in Figure 3, top-left panel), which indicated the direction of the next post. The arrows were necessary to indicate the intended facing direction when using the discordant teleporting interface and were present when using all locomotion interfaces. The arrow at the base of the green post pointed in the direction of the yellow post, the arrow at the base of the yellow post pointed in the direction of the red post, and

the arrow at the base of the red post pointed in the same direction as the arrow on the yellow post.

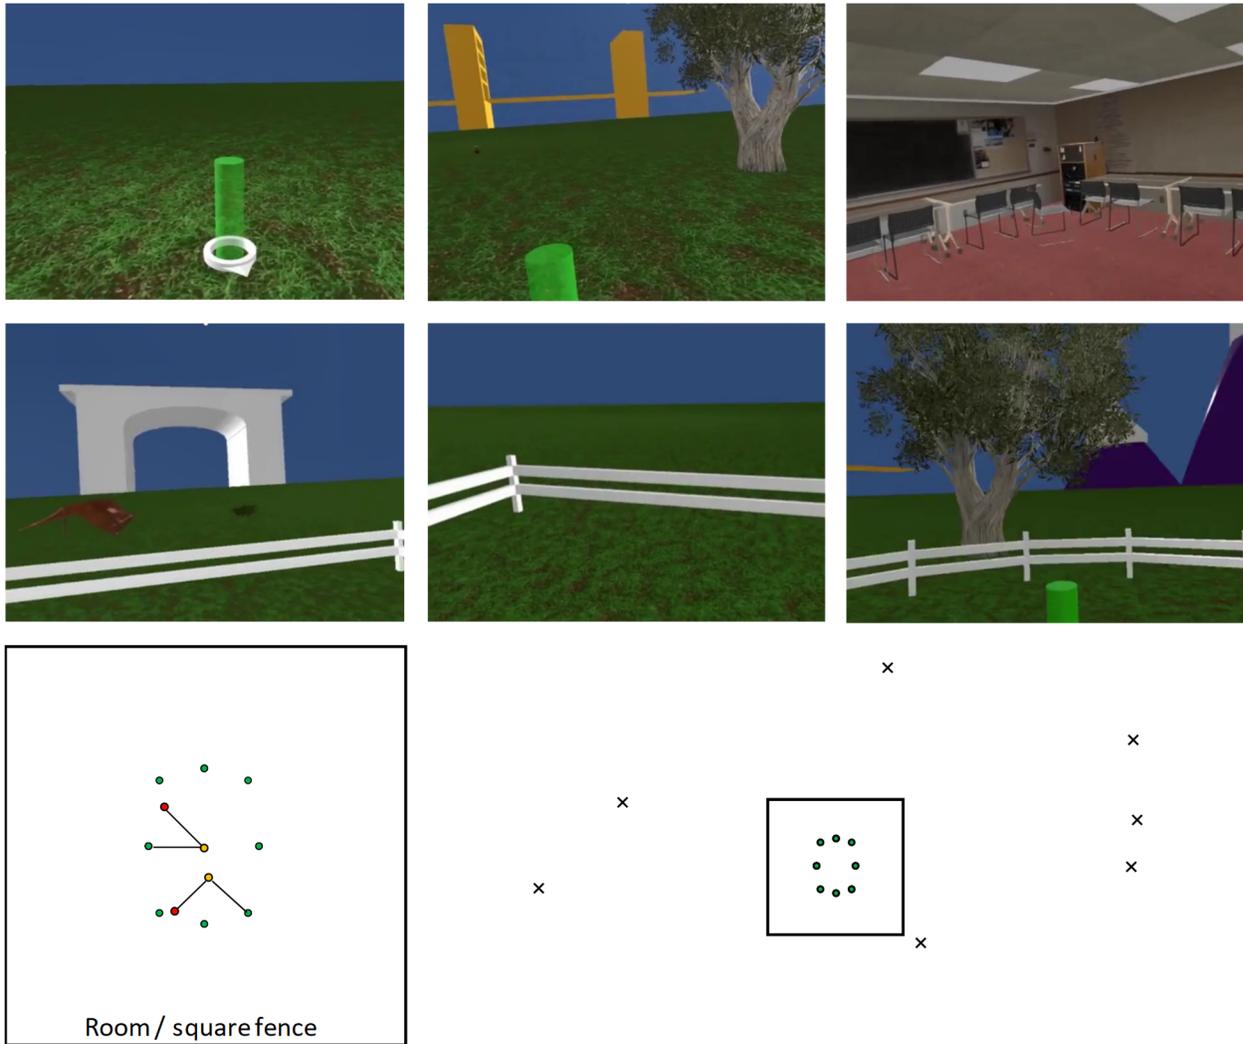


Figure 3. Virtual environments used in Experiments 1-5. Top row from left to right shows the open field (Exp 1, 2, and 3), landmarks (Exp 1), and classroom (Exp 2) environments. Middle row from left to right shows the square fence + landmarks (Exp 3 and 4), square fence only (Exp 4), and the circular fence + landmarks (Exp 5) environment. Bottom row shows overhead views of the VEs including possible post locations, boundaries, and landmarks (right panel only).

A virtual replica of the handheld controller was visible at all times, and its position and orientation were linked to that of the actual controller. The partially concordant teleporting interface was controlled by positioning a white circle (30 cm diameter) with surrounding white ring (75 cm diameter) in the intended location on the ground plane (see Figure 4, left panel). A thin red line extended from the joystick to the center of the white circle. The participant pressed and held the trackpad located on the top of the controller while manipulating the location of the teleport marker by pointing with the controller (similar to positioning a laser pointer). Releasing the trackpad teleported the participant to the selected location (orientation was unchanged). The completely discordant teleporting interface was controlled by positioning and orienting a magenta ring (height: 7.5 cm; outer diameter: 195 cm) with an arrow on one side (Figure 4, right panel). A thin red line extended from the joystick to the center of the ring. The participant pressed and held the trackpad button to bring up the teleporting ring, and rotated the ring by moving the thumb around the edge of the circular trackpad. Releasing the trackpad button teleported the participant to the selected location and orientation.

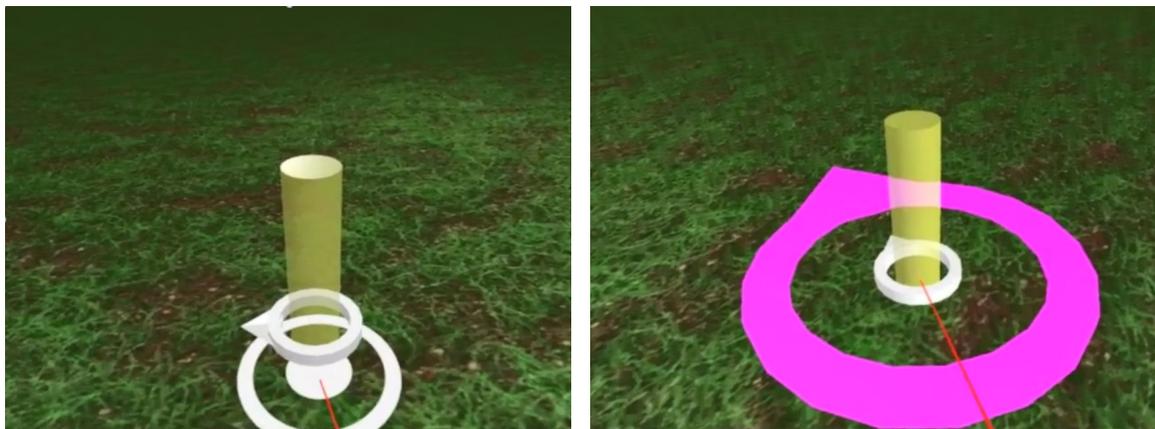


Figure 4. Screenshots taken from the participant's perspective while using the partially concordant teleporting interface (left panel) and the discordant teleporting interface (right panel).

Design.

The experiment employed a 2 (VE: open field or landmarks) by 3 (interface: walking, partially concordant teleporting, or discordant teleporting) repeated-measures design. For each combination of VE and interface, participants performed a block of 12 triangle completion trials corresponding to 12 unique turn angles spaced every 22.5° from -135° to $+135^\circ$. Path leg length was randomly selected on each trial from three possible values for the first and second leg (1.52, 1.68, or 1.83 m). There were eight possible locations of the path origin chosen to ensure that the participant would not navigate outside of the tracked space when using the walking interface. The first path leg led participants toward the center of the environment, and the second path leg led participants away from the center of the environment. Origin location was pseudo-randomized across trials with the constraint that the same origin location was not repeated on sequential trials.

The VE variable was blocked, such that three blocks corresponding to the three locomotion interfaces were completed first in one VE and then in the other VE. VE order was counterbalanced. Interface order was also counterbalanced, and the same order was used for both VE blocks. Pointing location and pointing response time were recorded.

Procedure.

After signing the informed consent, the participant was given verbal instructions on the triangle completion task. The participant then donned the HMD and was trained on performing the triangle completion task with each of the three locomotion interfaces. Training took place in a training VE with a grid-like texture on the ground plane and with no landmarks. The participant was required to perform three practice triangle completion trials with each

locomotion interface and could request additional practice. Experimental trials began after completion of practice.

A green post marking the path origin appeared at the beginning of each trial. The participant traveled to the green post using the assigned locomotion method. The green post disappeared upon arrival and a yellow post appeared, marking the first leg of the path. Upon arrival at the yellow post it disappeared and a red post appeared, marking the second leg of the path. Upon arrival at the red post, it disappeared, and the participant was prompted to point to the remembered location of the path origin. Pointing was accomplished by positioning a blue circle (38 cm diameter) on the ground plane. A thin red line extended from the joystick to the center of the blue circle. The participant pressed and held the trigger button located on the controller while manipulating the location of the blue circle and the experimenter then pressed a key to log the participant's response and advance to the next trial. Feedback was never provided.

Results

Absolute error was calculated as the absolute distance (in meters) between the point of origin and pointing response location. Response latency was calculated as the difference between the time when the participant arrived at the red post and when a pointing response was recorded. Data from one participant were removed due to incomplete data (at least one missing cell in the experimental design, attributable to experimenter error or failure to complete the study in the allotted time). Data from another participant were removed due to mean pointing errors that were more than three standard deviations higher than the group mean. An additional 38 trials (2.2%) were removed from the remaining data due to computer errors and procedural errors.

Analyses focused on the effects of interface and environment. Therefore, data from repeated trials for each environment and interface were averaged together prior to analysis. There

was no evidence of speed-accuracy tradeoff. The within-participant correlation between error and latency was significantly positive ($M = .47$, $SE = .09$), $t(23) = 5.29$, $p < .001$. Pointing error was the focus of the current project, and it was generally more responsive to manipulation of the independent variables than was response latency. Latency results are provided on the Open Science Framework: <https://osf.io/m4zfv/>.

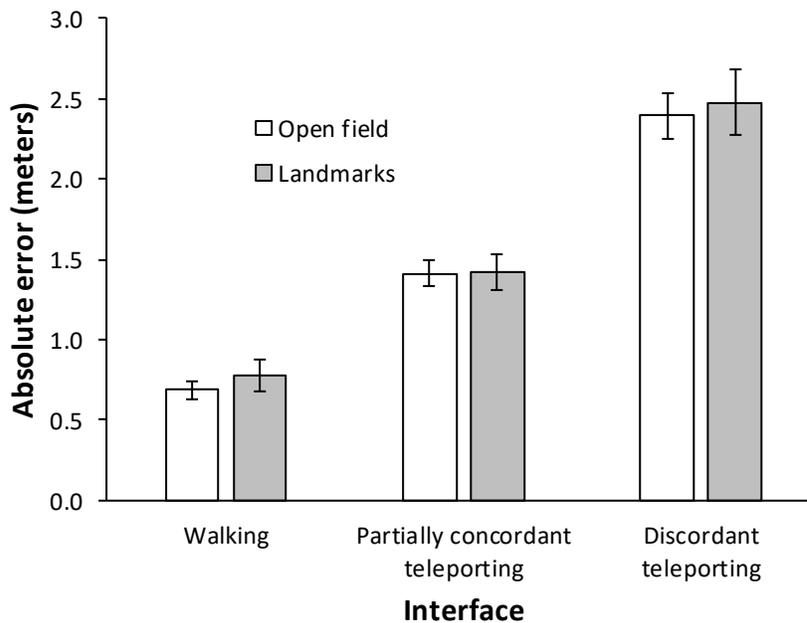


Figure 5. Average absolute error when performing the triangle completion task in Experiment 1. Error bars represent +/- 1 SEM.

Response errors (Figure 5) were analyzed in a repeated measures ANOVA with terms for interface and environment. Only the main effect of interface was significant, $F(2,46) = 114.19$, $p < .001$, $\eta_p^2 = .83$. The main effect of environment was not significant $F(1,23) = .60$, $p = .45$, $\eta_p^2 = .03$, nor was the interaction between interface and environment, $F(2,46) = .17$, $p = .84$, $\eta_p^2 = .01$. Errors in the walking interface ($M = 0.73$, $SE = 0.07$) were significantly lower than those in the partially concordant teleporting interface ($M = 1.42$, $SE = 0.09$), $t(23) = 8.51$, $p < .001$, $d = 1.75$,

and errors in the partially concordant teleporting interface were significantly lower than those in the discordant teleporting interface ($M = 2.44$, $SE = 0.16$), $t(23) = 8.12$, $p < .001$, $d = 1.66$.

Discussion

Spatial updating performance was best with the walking interface, worst with the completely discordant teleporting interface, and in between with the partially concordant teleporting interface. These results are consistent with the notion that body-based cues play an important role in spatial updating: the partially concordant teleporting interface lacked translational body-based cues, and the completely discordant teleporting interface lacked translational and rotational body-based cues. This pattern of results seems logical, yet is somewhat inconsistent with past research on the role of body-based cues in spatial updating. Using a very similar task, Klatzky et al. (1998) reported that removal of translational body-based cues does not impair spatial updating, but that removal of rotational body-based cues does. One potentially important difference is that the teleporting interfaces lacked optic flow in addition to body-based cues, whereas the study by Klatzky et al. provided optic flow in all conditions. This difference and others are considered in more detail in the General Discussion.

The more surprising result is that landmarks did not improve spatial updating performance compared to performance in the open field VE. Landmarks were expected to improve spatial updating with all interfaces, but particularly with the teleporting interfaces due to the lack of body-based cues. Past work indicates that body-based cues are integrated with landmark cues when walking (Chen et al., 2017; Nardini et al., 2008). However, those studies included landmarks that were relatively close to the target location, thereby providing both positional and directional information. It is possible that the landmarks in Experiment 1 were sufficiently far away as to be useful only as directional cues, although this still does not explain

why they were not helpful at all. Another factor is that the HMD limits the horizontal field of view to 100°, thereby limiting the possibility of using a collection of several landmarks to triangulate position.

Experiment 2 used an entirely different VE, a virtual classroom, in an effort to identify whether a richer and more naturalistic VE would support spatial updating, especially when teleporting. Although the differences between the landmarks VE and the classroom VE are so numerous that they defy succinct description, the landmarks VE is considered again later in Experiment 3 in light of insights gained through the classroom VE in Experiment 2.

Experiment 2

The landmarks VE from Experiment 1 did not facilitate spatial updating compared to the open field VE, even when navigating with the more difficult teleporting interfaces. Therefore, Experiment 2 utilized an indoor virtual environment composed of landmarks as well as walls, which provide a salient geometric cue to self-location (Sjolund et al., 2018). Research on spatial memory points to an important role for room shape in defining the reference frame used to organize remembered object locations (Kelly & McNamara, 2008; Shelton & McNamara, 2001). Furthermore, animal neuroscience has identified neurons that respond primarily to boundaries (Hartley, Burgess, Lever, Cacucci, & O'Keefe, 2000; Lever, Burton, Jeewajee, O'Keefe, & Burgess, 2009), and it is believed that these “boundary vector cells” are key inputs into the animal's representation of self-location (O'Keefe & Burgess, 1996; Burgess, Jackson, Hartley, & O'Keefe, 2000).

Participants completed the same triangle completion task from Experiment 1 using the same three navigation interfaces: walking, partially concordant teleporting, and discordant teleporting. The VE was an endless grassy field or a classroom.

Methods

Participants.

Twenty-seven students (13 men, 14 women) at Iowa State University participated in exchange for course credit. Data from three participants were removed (see Results) leaving 24 total participants (13 men, 11 women) in the dataset.

Stimuli, design, and procedure.

The same virtual reality system was used. The open field VE was identical to that used in Experiment 1. The classroom VE (Figure 3, top-right) was designed based on a real classroom at Iowa State University. The walls of the 3D model were textured with photographs from the real classroom. The classroom VE was square with 11-meter sides, and included several 3D models of classroom furniture such as chairs, tables, and a classroom media console (furniture and other virtual objects were placed near the perimeter of the room to ensure that the participant did not collide with virtual objects while performing the triangle completion task).

The experiment employed a 2 (VE: open field or classroom) by 3 (interface: walking, partially concordant teleporting, or discordant teleporting) repeated-measures design. The experimental design (trial numbers, trial block structure, independent and dependent variables) otherwise followed the Experiment 1 design.

Results

Data from three participants were removed due to incomplete data (at least one missing cell in the experimental design; attributable to experimenter error or failure to complete the study in the allotted time). An additional 26 trials (1.5%) were removed from the remaining data due to computer errors and procedural errors.

There was no evidence of speed-accuracy tradeoff. The within-participant correlation between error and latency was significantly positive ($M = .47$, $SE = .09$), $t(23) = 5.30$, $p < .001$.

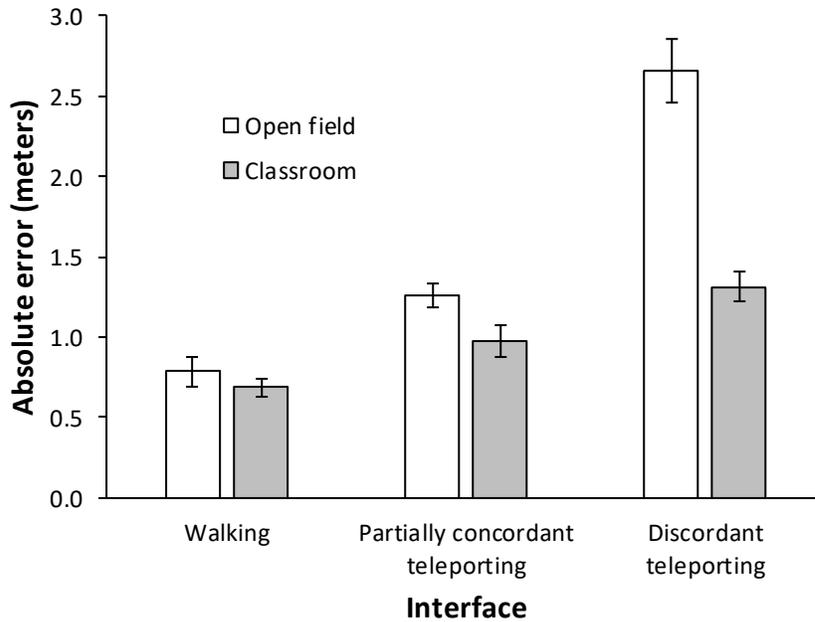


Figure 6. Average absolute error when performing the triangle completion task in Experiment 2. Error bars represent +/- 1 SEM.

Response errors (Figure 6) were analyzed in a repeated measures ANOVA with terms for interface and environment. Significant main effects of interface, $F(2,46) = 67.93$, $p < .001$, $\eta_p^2 = .75$, and environment, $F(1,23) = 38.31$, $p < .001$, $\eta_p^2 = .63$, were qualified by a significant interaction between interface and environment, $F(2,46) = 31.89$, $p < .001$, $\eta_p^2 = .58$. In each environment, the discordant teleporting interface produced larger errors than the partially concordant teleporting interface (open field: $t(23) = 7.08$, $p < .001$, $d = 1.44$; classroom: $t(23) = 3.65$, $p = .001$, $d = .74$), and the partially concordant teleporting interface produced larger errors than the walking interface (open field: $t(23) = 5.74$, $p < .001$, $d = 1.24$; classroom: $t(23) = 2.91$, $p = .008$, $d = .59$). Furthermore, the classroom produced lower errors compared to the open field

when using the partially concordant teleporting interface $t(23) = 2.37, p = .03, d = .48$, and the discordant teleporting interface $t(23) = 7.59, p < .001, d = 1.55$, but not when using the walking interface, $t(23) = 1.05, p = .30, d = .20$.

Discussion

Consistent with Experiment 1 results, spatial updating performance was best with the walking interface, worst with the completely discordant teleporting interface, and in between with the partially concordant teleporting interface. This pattern of results was found in the open field VE and the classroom VE. However, spatial updating performance in the classroom VE was better than in the open field VE when using the partially concordant and discordant teleporting interfaces. In other words, removal of body-based cues in the two teleporting interfaces negatively affected spatial updating in both VEs, but to a lesser extent in the classroom VE than in the open field VE. This result is in stark contrast to the landmarks VE in Experiment 1, which did not facilitate spatial updating performance with any of the three interfaces.

Why did the classroom VE, but not the landmarks VE, reduce spatial updating errors when using the two teleporting interfaces? Both VEs contained seemingly ample landmarks, although the landmarks differed in several ways, including proximity to the navigation space. One qualitative distinction between the VEs is that the classroom VE contained a boundary formed by four walls whereas the landmarks VE had no such boundary. Geometric cues defined by room walls affect human spatial memory (Kelly & McNamara, 2008; Shelton & McNamara, 2001) are integrated with self-motion cues during human spatial updating (Kelly et al., 2008; Sjolund et al., 2018). Furthermore, animal research has identified neurons that respond primarily to geometric boundaries (Lever et al., 2009).

Given the importance of spatial boundaries in navigation, Experiment 3 was designed to determine whether the landmarks VE from Experiment 1 would facilitate spatial updating if it also contained a geometric boundary.

Experiment 3

The classroom VE from Experiment 2 facilitated spatial updating compared to the open field VE when navigating with the teleporting interfaces, in contrast to the landmarks VE from Experiment 1, which showed no such facilitation. In order to examine the importance of boundaries, Experiment 3 added a square fence to the landmarks VE used in Experiment 1. The fence dimensions were the same as the classroom dimensions, and the fence style was chosen to allow a clear view of the landmarks beyond the fence borders. Spatial updating performance in this “fence + landmarks VE” was compared with performance in the open field VE used in Experiments 1 and 2.

Methods

Participants.

Twenty-eight students (11 men, 17 women) at Iowa State University participated in exchange for course credit. Data from four participants were removed (see Results) leaving 24 total participants (11 men, 13 women) in the dataset.

Stimuli, design, and procedure.

The same virtual reality system was used. The open field VE was identical to that used in Experiment 1. The fence + landmarks VE (Figure 3, middle-left) was identical to the landmarks VE from Experiment 1 except for the addition of a large square fence. Landmark locations were unchanged from Experiment 1, and all landmarks fell outside of the fence boundaries. The fence

was a split rail fence, consisting of two horizontal rails and four vertical posts, which allowed for visibility of landmarks beyond the fence.

The experiment employed a 2 (VE: open field or fence + landmarks) by 3 (interface: walking, partially concordant teleporting, or discordant teleporting) repeated-measures design. The experimental design (trial numbers, trial block structure, independent and dependent variables) otherwise followed the design of the prior experiments.

Results

Data from four participants were removed due to incomplete data (at least one missing cell in the experimental design; attributable to experimenter error or failure to complete the study in the allotted time). An additional 35 trials (2.0%) were removed from the remaining data due to computer errors and procedural errors.

There was no evidence of speed-accuracy tradeoff. The within-participant correlation between error and latency was significantly positive ($M = .50$, $SE = .07$), $t(23) = 6.86$, $p < .001$. Response errors (Figure 7) were analyzed in a repeated measures ANOVA with terms for interface and environment. Significant main effects of interface, $F(2,46) = 118.21$, $p < .001$, $\eta_p^2 = .84$, and environment, $F(1,23) = 33.72$, $p < .001$, $\eta_p^2 = .60$, were qualified by a significant interaction between environment and interface, $F(2,46) = 8.87$, $p = .001$, $\eta_p^2 = .28$. In each environment, the discordant teleporting interface produced larger errors than the partially concordant teleporting interface (open field: $t(23) = 6.05$, $p < .001$, $d = 1.24$; fence + landmarks: $t(23) = 4.31$, $p < .001$, $d = .89$), and the partially concordant teleporting interface produced larger errors than the walking interface (open field: $t(23) = 6.11$, $p < .001$, $d = 1.23$; fence + landmarks: $t(23) = 8.70$, $p < .001$, $d = 1.78$). Furthermore, the fence + landmarks VE produced lower errors compared to the open field when using the partially concordant teleporting interface, $t(23) =$

2.29, $p = .03$, $d = .45$, and the discordant teleporting interface, $t(23) = 4.26$, $p < .001$, $d = .87$, but not when using the concordant interface, $t(23) = 1.22$, $p = .24$, $d = .26$.

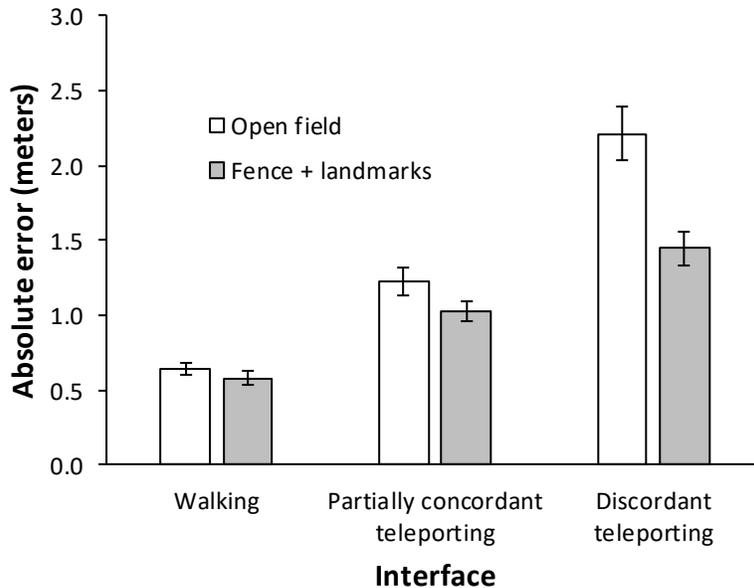


Figure 7. Average absolute error when performing the triangle completion task in Experiment 3.

Error bars represent +/- 1 SEM.

Follow-up testing was conducted to compare performance in the field + landmarks VE (Experiment 3) and the classroom VE (Experiment 2). Response errors were analyzed in a mixed model ANOVA with terms for interface and environment. Only the main effect of interface was significant, $F(2,92) = 65.68$, $p < .001$, $\eta_p^2 = .59$. The main effect of environment was not significant, $F(1,46) = 0.08$, $p = .78$, nor was the interaction, $F(2,92) = 1.70$, $p = .19$.

Discussion

Consistent with results from Experiments 1 and 2 spatial updating performance was best with the walking interface, worst with the completely discordant teleporting interface, and in between with the partially concordant teleporting interface. Although this pattern of results was

found in both the open field VE and the fence + landmarks VE, spatial updating performance in the fence + landmarks VE was better than in the open field VE when using the partially concordant and discordant teleporting interfaces. This benefit of the fence + landmarks VE when using the two teleporting interfaces mirrors the performance benefit found in the classroom VE in Experiment 2. Furthermore, spatial updating performance did not differ between the fence + landmarks VE in Experiment 3 and the classroom VE in Experiment 2.

The facilitative influence of the fence + landmarks VE stands in stark contrast to the landmarks VE in Experiment 1, which did not facilitate spatial updating performance with any of the three interfaces. Instead, performance in the fence + landmarks VE was comparable to that in the classroom VE. Therefore, the presence of a rectangular boundary may be the critical distinction between the landmarks VE in Experiment 1 and the classroom VE in Experiment 2. Because the fence + landmarks VE contained a rectangular boundary and several landmarks it is impossible to say whether the boundary is sufficient to facilitate spatial updating performance when teleporting or whether landmarks are also a necessary component. Experiment 4 explores the sufficiency of a square boundary.

Although it is possible that participants did not use the landmarks and instead relied exclusively on the fence in the fence + landmarks VE, it is also possible that the fence led participants to use the landmarks. The square fence by itself is somewhat ambiguous since each side is visually identical. Therefore, the landmarks may have been useful for providing an orientation to the boundary. That is, the landmarks can be used to identify the “north” side of the boundary. Although this may not be necessary in a square enclosure because the square itself might be sufficient to keep track of “north,” a more ambiguous boundary could be used to identify whether landmarks are especially useful as tools to provide orientation to boundaries.

Therefore, Experiment 5 explores whether landmarks can be useful for disambiguating a circular boundary.

Experiment 4

The fence + landmarks VE from Experiment 3 facilitated spatial updating compared to the open field VE when navigating with the teleporting interfaces, in contrast to the landmarks VE from Experiment 1. In order to examine the sufficiency of boundaries, Experiment 4 evaluated spatial updating performance in a “fence only VE,” which contained the same fence used in the fence + landmarks VE but without the landmarks. Performance in the fence only VE was compared to the same fence + landmarks VE used in Experiment 3.

Methods

Participants.

Twenty-nine students (15 men, 14 women) at Iowa State University participated in exchange for course credit. Data from five participants were removed (see Results) leaving 24 total participants (12 men, 12 women) in the dataset.

Stimuli, design, and procedure.

The same virtual reality system was used. The fence + landmarks VE was identical to that used in Experiment 3. The fence only VE was created by removing all of the landmarks from the fence + landmarks VE, leaving only the fence and grassy field.

The experiment employed a 2 (VE: fence only or fence + landmarks) by 3 (interface: walking, partially concordant teleporting, or discordant teleporting) repeated-measures design. The experimental design (trial numbers, trial block structure, independent and dependent variables) otherwise followed the design of the prior experiments.

Results

Data from five participants were removed due to incomplete data (at least one missing cell in the experimental design; attributable to experimenter error or failure to complete the study in the allotted time). An additional 33 trials (1.91%) were removed from the remaining data due to computer errors and procedural errors.

There was no evidence of speed-accuracy tradeoff. The within-participant correlation between error and latency was significantly positive ($M = .57$, $SE = .06$), $t(23) = 9.89$, $p < .001$. Response errors (Figure 8) were analyzed in a repeated measures ANOVA with terms for interface and environment. The main effect of interface was significant, $F(2,46) = 48.96$, $p < .001$, $\eta_p^2 = .68$, as was the interaction between environment and interface, $F(2,46) = 4.72$, $p = .01$, $\eta_p^2 = .17$. The main effect of environment was not significant, $F(1,23) = 2.02$, $p = .17$. In each environment, the discordant teleporting interface produced larger errors than the partially concordant teleporting interface (fence only: $t(23) = 6.22$, $p < .001$, $d = 1.27$; fence + landmarks: $t(23) = 4.25$, $p < .001$, $d = .88$), and the partially concordant teleporting interface produced larger errors than the walking interface (fence only: $t(23) = 5.37$, $p < .001$, $d = 1.08$; fence + landmarks: $t(23) = 4.60$, $p < .001$, $d = .93$). Furthermore, the fence + landmarks VE produced lower errors compared to the fence only VE when using the discordant teleporting interface $t(23) = 2.17$, $p = .04$, $d = .43$, but not when using the partially concordant teleporting interface and the walking interface (partially concordant teleporting: $t(23) = .33$, $p = .75$, $d = .07$; walking: $t(23) = .81$, $p = .43$, $d = .16$).

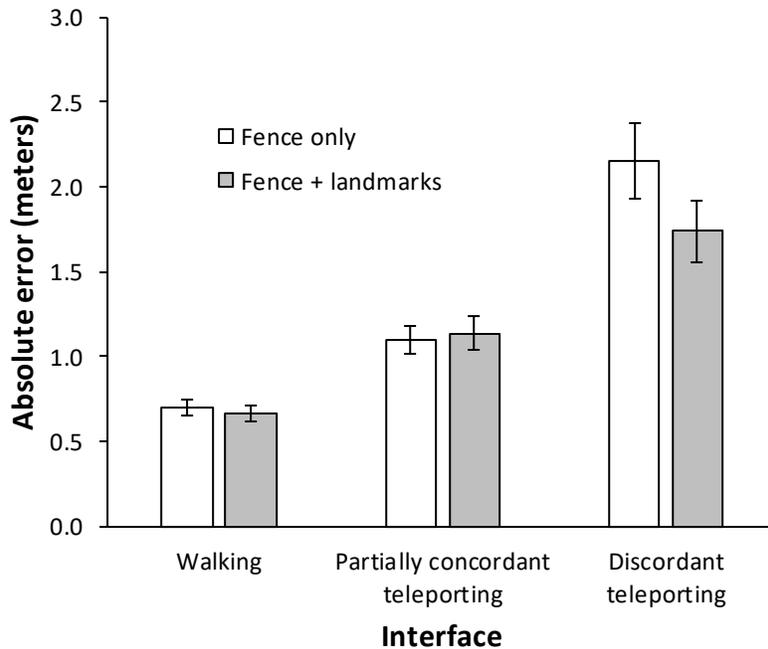


Figure 8. Average absolute error when performing the triangle completion task in Experiment 4.

Error bars represent ± 1 SEM.

Discussion

Consistent with results from Experiments 1–3, spatial updating performance was best with the walking interface, worst with the completely discordant teleporting interface, and in between with the partially concordant teleporting interface. Although this pattern of results was found in both the fence + landmarks VE and the fence only VE, spatial updating performance in the fence + landmarks VE was better than in the open field VE when using the discordant teleporting interface. These results indicate that a boundary without landmarks can be sufficient and that additional landmarks are not necessary when all body-based cues are available (walking interface) or when only rotational body-based cues are available (partially concordant teleporting interface), but that landmarks visible beyond the boundary improve spatial updating performance when no body-based cues are available (discordant teleporting interface).

Why were landmarks (in addition to the boundary) helpful only when using the discordant teleporting interface? The square boundary is a visual cue to self-orientation, but it is somewhat ambiguous due to its rotational symmetry. The visual boundary is an informative cue when the navigator can estimate self-orientation within $\pm 45^\circ$, but if uncertainty exceeds that threshold, the navigator could confuse two adjacent sides of the boundary. Therefore, one explanation for why landmarks were helpful only with the discordant teleporting interface is that participants were generally able to estimate self-orientation within $\pm 45^\circ$ with the partially concordant teleporting interface, but at least occasionally exceeded that threshold with the discordant teleporting interface. One implication for future research is that an unambiguous boundary should lead to equivalent spatial updating performance whether or not landmarks are included.

Experiment 5

The fence + landmarks VE from Experiment 3 facilitated spatial updating compared to the open field VE when navigating with the teleporting interfaces, in contrast to the landmarks VE from Experiment 1. We speculate that the value of landmarks in this context is that they define a boundary's orientation. If true, then even a circular boundary should benefit from added landmarks beyond the boundary. This would provide particularly strong support of the theory since landmarks alone (Experiment 1) did not facilitate spatial updating, and a circular fence alone should also not be especially useful. Experiment 5 compared spatial updating performance with a circular boundary only ("circle only VE") and a circular boundary with the same landmarks used in Experiment 1 ("circle + landmarks VE").

Methods

Participants.

Twenty-eight students (10 men, 18 women) at Iowa State University participated in exchange for course credit. Data from four participants were removed (see Results) leaving 24 total participants (9 men, 15 women) in the dataset.

Stimuli, design, and procedure.

The same virtual reality system was used. The circle + landmarks VE (Figure 3, middle-right) used the same proximal and distal landmarks from the landmarks VE condition in Experiment 1 and added a circular fence (11.3 m diameter). The circle only VE was created by removing all of the landmarks from the circle + landmarks VE, leaving only the circular fence and grassy field.

The experiment employed a 2 (VE: circle only or circle + landmarks) by 3 (interface: walking, partially concordant teleporting, or discordant teleporting) repeated-measures design. The experimental design (trial numbers, trial block structure, independent and dependent variables) otherwise followed the design of the prior experiments.

Results

Data from four participants were removed due to incomplete data (at least one missing cell in the experimental design; attributable to experimenter error or failure to complete the study in the allotted time). An additional 39 trials (2.26%) were removed from the remaining data due to computer errors and procedural errors.

There was no evidence of speed-accuracy tradeoff. The within-participant correlation between error and latency was significantly positive ($M = .60$, $SE = .07$), $t(23) = 8.08$, $p < .001$.

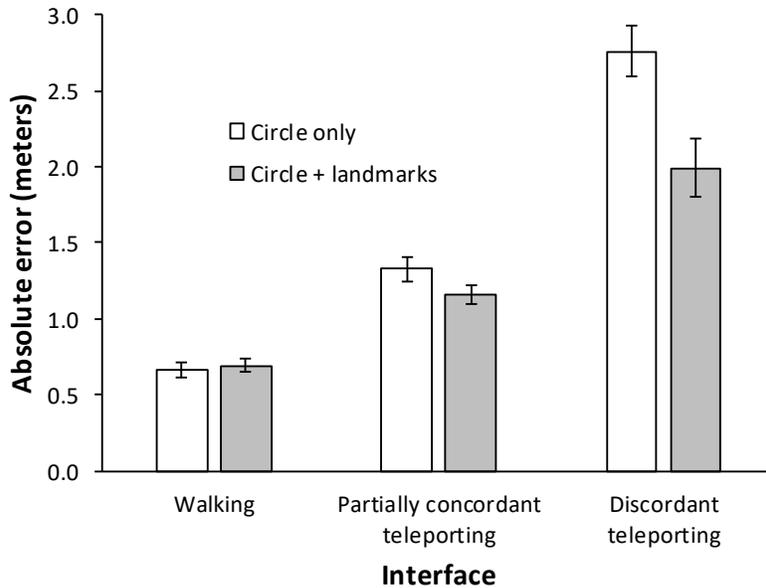


Figure 9. Average absolute error when performing the triangle completion task in Experiment 5.

Error bars represent +/- 1 SEM.

Response errors (Figure 9) were analyzed in a mixed model ANOVA with terms for interface and environment. Significant main effects of interface, $F(2,46) = 121.07, p < .001, \eta_p^2 = .84$, and environment $F(1,23) = 12.62, p = .002, \eta_p^2 = .35$, were qualified by a significant interaction between environment and interface, $F(2,46) = 7.84, p = .001, \eta_p^2 = .25$. In each environment, the discordant teleporting interface produced larger errors than the partially concordant teleporting interface (circle only: $t(23) = 7.85, p < .001, d = 1.61$; circle + landmarks: $t(23) = 4.70, p < .001, d = .95$), and the partially concordant teleporting interface produced larger errors than the walking interface (circle only: $t(23) = 7.66, p < .001, d = 1.57$; circle + landmarks: $t(23) = 6.38, p < .001, d = 1.31$). The circle + landmarks VE produced lower errors compared to the circular fence when using the discordant teleporting interface, $t(23) = 3.29, p = .003, d = .67$, and the same pattern was found when using the partially concordant teleporting interface but the

difference did not reach statistical significance, $t(23) = 1.80$, $p = .08$, $d = .36$. There was no benefit to the added landmarks when using the walking interface, $t(23) = .65$, $p = .52$, $d = .14$.

Discussion

Compared to the circle only condition in which the only environmental cue was a circular fence, the circle + landmarks condition led to better spatial updating performance when using the discordant teleporting interface and the partially concordant teleporting interface, although the latter difference was marginally significant. These data support the theory that landmarks are especially valuable when they can be used to provide orientation to a boundary and stand in contrast to the finding from Experiment 1 that landmarks without a boundary do not improve performance.

General Discussion

The primary goals of this project were to evaluate whether there are spatial cognitive costs associated with the use of two teleporting interfaces commonly found in virtual reality applications and to identify environmental cues that could mitigate those costs. Participants completed a triangle completion task using three locomotion interfaces: 1) walking, which included body-based and visual self-motion cues, 2) partially concordant teleporting, which included body-based and visual rotational cues but no translational cues, and 3) discordant teleporting, which lacked all self-motion cues. Across five experiments, discordant teleporting produced larger errors than partially concordant teleporting, which in turn produced larger errors than walking. Furthermore, geometric boundaries (room walls or a fence) were necessary to mitigate the spatial cognitive costs associated with teleporting, and landmarks alone were unhelpful.

Based on past work showing that body-based rotational cues are both necessary and sufficient for triangle completion (Klatzky et al., 1998), it was predicted that walking and partially concordant teleporting would produce equally low triangle completion errors compared to discordant teleporting. However, the actual pattern observed across the five experiments reported here is only somewhat consistent with the results of Klatzky et al. (1998), with the main departure being the larger errors with the partially concordant teleporting interface compared to the walking interface. However, there were several methodological differences across studies. For example, the partially concordant and discordant teleporting interfaces in the current study did not provide translational optic flow, all conditions involved active selection of locomotion destination, and all conditions provided information about the turn angle prior to executing the turn (i.e., the arrow on the yellow post pointed toward the unseen red post). Of these differences, only the absence of translational optic flow seems potentially detrimental to performance, and therefore seems the most likely cause of increased error with partially concordant teleporting compared to walking. A future experiment in which the partially concordant interface is tested with and without translational optic flow could help answer this question, but it is not practical in an applied context due to the increased sickness when translational optic flow is provided in the absence of translational body movement (Christou & Aristidou, 2017; Moghadam et al., 2018; Langbehn et al., 2018; Weißker et al., 2018).

It was also predicted that the availability of piloting cues would result in lower errors for all locomotion interfaces, and especially so for discordant teleporting. Surprisingly, a collection of proximal and distal landmarks (Experiment 1) did not reduce triangle completion errors with any of the three locomotion interfaces. A virtual classroom (Experiment 2) replete with landmark and geometric boundary cues did facilitate performance with the two teleporting interfaces. The

geometric boundary seems to be critical, as adding a square fence to the landmarks VE (Experiment 3) facilitated performance with the two teleporting interfaces in a manner indistinguishable from the classroom VE. The square boundary in the absence of landmarks (Experiment 4) was sufficient to improve performance for the partially concordant teleporting interface, but not the discordant teleporting interface. Further highlighting the importance of boundaries, adding a circular fence to the landmarks VE (Experiment 5) was sufficient to improve performance with the teleporting interfaces. Collectively, these results indicate that geometric boundaries can improve spatial updating when using a teleporting interface and that landmarks can be helpful but only in the context of a geometric boundary.

The experiments that included outdoor landmarks indicate that landmarks are helpful primarily to provide orientation to a boundary. Whereas landmarks alone did not facilitate spatial updating (Experiment 1), landmarks in combination with a surrounding fence did facilitate spatial updating. The experiment with the circular fence (Experiment 5) best illustrates the value of landmarks for providing orientation to a boundary. The circular fence alone was essentially useless for determining self-orientation during travel due to its rotational symmetry, and triangle completion performance in that environment was no different from the open field environment from other experiments¹. However, the landmark cues presented beyond the fence border provided an orientation to the boundary cue that allowed for more accurate identification of self-orientation.

¹ A 2×3 mixed-model ANOVA with terms for environment (circular fence from Experiment 5 or open field from Experiments 1-3) and interface showed only a significant main effect of interface, $F(2,188) = 268.07, p < .001, \eta_p^2 = .74$. The main effect of environment was not significant, $F(1,94) = 1.405, p = .239, \eta_p^2 = .02$. The interaction between environment and interface was marginal, $F(2,188) = 2.88, p = .059, \eta_p^2 = .03$, with errors slightly higher in the discordant interface in the circular fence environment compared to the open field.

Even when piloting cues improved triangle completion performance with the teleporting interfaces, those same piloting cues did not improve performance when walking. This result seems at odds with research showing that body-based locomotion cues are integrated with piloting cues to improve spatial updating performance (Chen et al., 2017; Sjolund et al., 2018; Nardini et al., 2008). However, cue combination research also shows that multiple cues are weighted based on their predictive value, and if one cue is much more predictive than another, it is difficult to distinguish whether the cues are being combined or if the most predictive cue is being followed exclusively. Since no attempt was made to equate the predictive value of body-based walking cues and visual piloting cues, it is possible that one cue (probably the body-based cues) simply dominated when all cues were available.

Given the importance of boundaries for accurate spatial updating, future research should investigate whether the form of the boundary affects its value as a cue to spatial updating. Architects have long considered that the design of spaces and subspaces can influence how people feel and behave, and the boundaries of those spaces can be physical (land, water, railroads), or simply “in people’s minds” (Alexander et al., 1977). Boundary vector cells in the rodent brain respond to environmental boundaries such as walls or drop-offs (Lever et al., 2009). In humans, the representation of purely visual boundaries such as flooring transitions appears to be dissociable from navigational boundaries (Julian, Ryan, Hamilton, & Epstein, 2016). It would be interesting to evaluate whether human spatial updating is primarily affected by navigationally-relevant boundaries or whether visual boundaries are equally useful.

The teleporting interface is widely used in VR applications, and its advantages include ease of use (Bozgeyikli et al., 2016; Langbehn et al., 2018) and low incidence of sickness (Christou & Aristidou, 2017; Moghadam et al., 2018; Langbehn et al., 2018; Weißker et al.,

2018). However, the experiments reported here indicate that developers should be cautious about the potential for disorientation, especially when users teleport to change their location and orientation. When teleportation is used, environmental boundaries can mitigate but not eliminate the spatial cognitive consequences of teleporting, and landmarks can be helpful when presented in the context of boundaries.

References

- Alexander, C., Ishikawa, S., Silverstein, M., Jacobson, M., Fiksdahl-King, I., & Angel, S. (1977). *A Pattern Language: Towns, Buildings, Construction*. New York: Oxford University Press.
- Bowman, D. A., Koller, D., & Hodges, L. F. (1997). Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. *Proceedings of the 1997 Virtual Reality Annual International Symposium*, 45-52.
- Bozgeyikli, E., Raij, A., Katkooori, S., & Dubey, R. (2016). Point & teleport locomotion technique for virtual reality. *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play*, 205-216. <https://doi.org/10.1145/2967934.2968105>
- Burgess, N., Jackson, A., Hartley, T., & O'Keefe, J. (2000). Predictions derived from modeling the hippocampal role in navigation. *Biological Cybernetics*, 83, 301-312.
- Chance, S. S., Gaunet, F., Beall, A. C., & 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.
- Chen, X., McNamara, T. P., Kelly, J. W., & Wolbers, T. (2017). Cue combination in human spatial navigation. *Cognitive Psychology*, 95, 105-144.
<https://doi.org/10.1016/j.cogpsych.2017.04.003>
- Chrastil, E. R. & Waren, W. H. (2012). Active and passive contributions to spatial learning. *Psychonomic Bulletin & Review*, 19, 1-23.
- Christou, C. G. & Aristidou, P. (2017). Steering Versus Teleport Locomotion for Head Mounted Displays. In De Paolis, L., Bourdot, P., & Mongelli, A. (Eds), *Augmented Reality, Virtual*

- Reality, and Computer Graphics: Lecture Notes in Computer Science, vol 10325* (pp. 431-446). Springer.
- Grant, S. C., & Magee, L. E. (1998). Contributions of proprioception to navigation in virtual environments. *Human Factors, 40*(3), 489-497.
- Grechkin, T., Thomas, J., Azmandian, M., Bolas, M. & Suma, E. (2016). Revisiting detection thresholds for redirected walking: Combining translation and curvature gains. *Proceedings of the ACM Symposium on Applied Perception*, 113–120.
- Hartley, T., Burgess, N., Lever, C., Cacucci, F., & O’Keefe, J. (2000). Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus, 10*, 369-379.
- Interrante, V., Ries, B., & Anderson, L. (2007). Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments. *3DUI '07: Proceedings of the 2007 IEEE Symposium on 3D User Interfaces*, 167-170.
- Julian, J. B., Ryan, J., Hamilton, R. H., & Epstein, R. A. (2016). The occipital place area is causally involved in representing environmental boundaries during navigation. *Current Biology, 26*(8), 1104-1109. <https://doi.org/10.1016/j.cub.2016.02.066>.
- Kearns, M. J., Warren, W. H., Duchon, A. P., & Tarr, M. J. (2002). Path integration from optic flow and body senses in a homing task. *Perception, 31*, 349-374. DOI:10.1068/p3311
- Kelly, J.W. & McNamara, T.P. (2008). Response mode differences in perspective taking: Differences in representation or differences in retrieval? *Memory & Cognition, 36*(4), 863-872.
- Kelly, J. W., McNamara, T. P., Bodenheimer, B., Carr, T. H., & Rieser, J. J. (2008). The shape of human navigation: How environmental geometry is used in maintenance of spatial orientation. *Cognition, 109*(2), 281-286. doi:10.1016/j.cognition.2008.09.001

- Klatzky, R. L., Loomis, J. M., Beall, A. C., Chance, S. S., & Golledge, R. G. (1998). Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science, 9*, 293-298.
- Langbehn, E., Lubos, P., & Steincke, F. (2018). Evaluation of locomotion techniques for room-scale VR: Joystick, teleportation, and redirected walking. *Proceedings of the Virtual Reality International Conference 2018*, 1-9. <https://doi.org/10.1145/3234253.3234291>.
- Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *The Journal of Neuroscience, 29*(31), 9771-9777. DOI:10.1523/JNEUROSCI.1319-09.2009
- Mainelli, T., Shirer, M. & Ubrani, J. (2019). Augmented reality and virtual reality headsets poised for significant growth. International Data Corporation. Retrieved from <https://www.idc.com/getdoc.jsp?containerId=prUS44966319>
- May, M. (2004). Imaginal perspective switches in remembered environments: Transformations versus interface accounts. *Cognitive Psychology, 48*, 163-206. doi:10.1016/S0010-0285(03)00127-0
- Moghadam, K. R., Banigan, C., & Ragan, E. D. (2018). Scene transitions and teleportation in virtual reality and the implications for spatial awareness and sickness. *IEEE Transactions on Visualization and Computer Graphics* <https://doi.org/10.1109/TVCG.2018.2884468>
- Nardini, M., Jones, P., Bedford, R., & Braddick, O. (2008). Development of cue integration in human navigation. *Current Biology, 18*(9), 689-693. <https://doi.org/10.1016/j.cub.2008.04.021>
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature, 381*, 425-428.

- Padilla, L. M., Creem-Regehr, S. H., Stefanucci, J. K., & Cashdan, E. A. (2017). Sex differences in virtual navigation influenced by scale and navigation experience. *Psychonomic Bulletin & Review*, *24*(2), 582-590. DOI 10.3758/s13423-016-1118-2
- Presson, C. C., & Montello, D. R. (1994). Updating after rotational and translational body movements: Coordinate structure of perspective space. *Perception*, *23*, 1447-1455.
- Riecke, B. E., van Veen, H. A. H. C., & Bühlhoff, H. H. (2002). Visual homing is possible without landmarks: A path integration study in virtual reality. *Presence Teleoperators and Virtual Environments*, *11*(5), 443-473.
- Rieser, J. J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(6), 1157-1165.
- Ruddle, R. A., & Lessels, S. (2006). For efficient navigational search, humans require full physical movement, but not a rich visual scene. *Psychological Science*, *17*(6), 460-465.
- Ruddle, R. A., Volkova, E., & Bühlhoff, 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*(2), 1-22. <https://doi.org/10.1145/1970378.1970384>
- Ruddle, R. A., Volkova, E., Mohler, B., & Bühlhoff, H. H. (2011). The effect of landmark and body-based sensory information on route knowledge. *Memory & Cognition*, *39*(4), 686-699. <https://doi.org/10.3758/s13421-010-0054-z>
- Shelton, A. L., & McNamara, T. P. (2001). Systems of spatial reference in human memory. *Cognitive Psychology*, *43*, 274-310. doi:10.1006/cogp.2001.0758

- Sjolund, L. A., Kelly, J. W., & McNamara, T. P. (2018). Optimal combination of environmental cues and path integration during navigation. *Memory & Cognition*, *46*(1), 89-99.
<https://doi.org/10.3758/s13421-017-0747-7>
- Steinicke, F., Bruder, G., Jerald, J., Frenz, H., & Lappe, M. (2010). Estimation of detection thresholds for redirected walking techniques. *IEEE Transactions on Visualization and Computer Graphics*, *16*(1), 17–27.
- Sun, H., Chan, G. S. W., & Campos, J. L. (2004). Active navigation and orientation-free spatial representations. *Memory & Cognition*, *32*, 51–71.
- Waller, D., Loomis, J. M., Golledge, R. G., & Beall, A. C. (2000). Place learning in humans: The role of distance and direction information. *Spatial Cognition and Computation*, *2*, 333-354.
- Waller, D., Loomis, J. M., & Haun, D. B. (2004). Body-based senses enhance knowledge of directions in large-scale environments. *Psychonomic Bulletin & Review*, *11*(1), 157-163.
- Weißker, T., Kunert, A., Fröhlich, B., & Kulik, A. (2018). Spatial updating and simulator sickness during steering and jumping in immersive virtual environments. *2018 IEEE Conference on Virtual Reality and 3d User Interfaces (VR)*, 97-104.
<https://doi.org/10.1109/VR.2018.8446620>
- Williams, B., Narasimham, G., McNamara, T. P., Carr, T. H., Rieser, J. J., & Bodenheimer, B. (2006). Updating orientation in large virtual environments using scaled translational gain. *Proceedings of the 3rd Symposium on Applied Perception in Graphics and Visualization*, 21–28.
- Wraga, M., Creem-Regehr, S. H., & Proffitt, D. R. (2004). Spatial updating of virtual displays during self- and display rotation. *Memory & Cognition*, *32*(3), 399-415.