Summer 2013

Cognitive Demands of Semi-Natural Virtual Locomotion

William E. Marsh
Iowa State University

Jonathan W. Kelly
Iowa State University, jonkelly@iastate.edu

Veronica J. Dark
Iowa State University, vjdark@iastate.edu

James H. Oliver
Iowa State University, oliver@iastate.edu

The complete bibliographic information for this item can be found at http://lib.dr.iastate.edu/me_pubs/52. For information on how to cite this item, please visit http://lib.dr.iastate.edu/howtocite.html.

Cognitive Demands of Semi-Natural Virtual Locomotion

Abstract

There is currently no fully natural, general-purpose locomotion interface. Instead, interfaces such as gamepads or treadmills are required to explore large virtual environments (VEs). Furthermore, sensory feedback that would normally be used in real-world movement is often restricted in VR due to constraints such as reduced field of view (FOV). Accommodating these limitations with locomotion interfaces afforded by most virtual reality (VR) systems may induce cognitive demands on the user that are unrelated to the primary task to be performed in the VE. Users of VR systems often have many competing task demands, and additional cognitive demands during locomotion must compete for finite resources. Two studies were previously reported investigating the working memory demands imposed by semi-natural locomotion interfaces (Study 1) and reduced sensory feedback (Study 2). This paper expands on the previously reported results and adds discussion linking the two studies. The results indicated that locomotion with a less natural interface increases spatial working memory demands, and that locomotion with a lower FOV increases general attentional demands. These findings are discussed in terms of their practical implications for selection of locomotion interfaces when designing VEs.

Keywords

VRAC, Psychology

Disciplines

Cognition and Perception | Mechanical Engineering

Comments

This article is from *Presence: Teleoperators and Virtual Environments* 22 (2013): 216–234, doi:10.1162/PRES_a_00152. Posted with permission.

William E. Marsh*
Jonathan W. Kelly
Veronica J. Dark
James H. Oliver
Iowa State University
Virtual Reality Applications Center
Ames, IA 50011

Cognitive Demands of Semi-Natural Virtual Locomotion

Abstract

There is currently no fully natural, general-purpose locomotion interface. Instead, interfaces such as gamepads or treadmills are required to explore large virtual environments (VEs). Furthermore, sensory feedback that would normally be used in real-world movement is often restricted in VR due to constraints such as reduced field of view (FOV). Accommodating these limitations with locomotion interfaces afforded by most virtual reality (VR) systems may induce cognitive demands on the user that are unrelated to the primary task to be performed in the VE. Users of VR systems often have many competing task demands, and additional cognitive demands during locomotion must compete for finite resources. Two studies were previously reported investigating the working memory demands imposed by semi-natural locomotion interfaces (Study 1) and reduced sensory feedback (Study 2). This paper expands on the previously reported results and adds discussion linking the two studies. The results indicated that locomotion with a less natural interface increases spatial working memory demands, and that locomotion with a lower FOV increases general attentional demands. These findings are discussed in terms of their practical implications for selection of locomotion interfaces when designing VEs.

1 Introduction

Virtual reality (VR) is often used in domains involving simultaneous objectives competing for a user's limited cognitive resources. For example, a VR firefighting simulation presents the user with multiple visual cues regarding the state of the fire and the burning structure and also requires the user to make a decision about the best course of action in light of those cues. Finally, the user must execute the selected action within the virtual environment.

Virtual environments (VEs) are often very large, but they must be manipulated from within the finite bounds of a VR system. Unlike the case for real environments, physically walking through the VE is often impossible due to space constraints, and alternate interfaces, such as joysticks or wands, must be substituted for physical walking. Interaction with these devices differs from real-world movements and may impact performance.

Many VEs are designed to simulate real-world scenarios, but the experience is constrained by the VR hardware used. For example, typical VR systems

Presence, *Vol. 22, No. 3, Summer 2013,* 216–234
doi:10.1162/PRES_a_00152

---

*Correspondence to william@wemarsh.com.

restrict the user's field of view, as compared to real-world interactions. This can decrease the feeling of presence in the virtual environment, and the lack of expected sensory feedback may hinder a user's virtual movement performance.

Because of these limitations in the control actions and in the sensory feedback resulting from movements, no general interface for virtual locomotion is truly natural. These unnatural aspects may cause users to employ strategies with additional cognitive demands. If so, these strategies may compete for cognitive resources from the same pools that are utilized for successful completion of a user's primary domain-related tasks.

1.1 Virtual Locomotion and Navigation

The type of movement required to move through space depends on the scale of that space. Montello (1993) proposed four scales of space: figural, vista, environmental, and geographical. Different scales have different movement requirements for exploration. Figural space is small relative to the body, and includes objects or pictures that can be manipulated with the hands. Vista space can mostly be seen from a given vantage point, and includes room-sized spaces. Figural and vista spaces can be apprehended in their entirety (or nearly so) from a single vantage point, so changing vantage points in those spaces may not be necessary. However, movement through the environment has the potential to provide additional cues about environmental layout (Gibson, 1966), so self-motion may be important even at figural and vista scales. Environmental space is large relative to the body, and includes buildings or even cities. Environmental spaces require navigation in order to experience the entire scene and these spaces are therefore a frequent focus of virtual locomotion research. Geographical space is even larger and cannot be explored with locomotion alone.

In this paper, the term navigation refers to activities requiring large-scale representations of space, such as wayfinding and path integration, which are not considered in detail here. The term virtual locomotion is used to refer to the atomic actions that are pieced together when a person navigates through environmental space, and a locomotion interface refers to a set of controls or movements required to effect those actions.

A typical adult has many years of experience interacting with the physical world, so one way to increase naturalness in VR is to mimic those interactions as closely as possible (Kulik, 2009; Wickens & Baker, 1995). A natural locomotion interface should also maximize the match between actions, proprioceptive information, and sensory feedback generated by the VR system, with respect to analogous actions in the physical world. A more natural interface is likely to place fewer physical and cognitive demands on the user. However, because VEs are often large, implementing a completely natural interface is not typically possible. Selecting an interface usually involves tradeoffs, and different interface choices may cause users to employ different cognitive strategies.

Many locomotion interfaces currently exist. The gamepad is a commonly used interface because it is easy to implement, inexpensive, and allows for infinite virtual locomotion from within constrained physical spaces. These devices typically have two sticks: one for rotation and one for translation velocity in the VE. The gamepad is considered to be a rate-control device because user input is translated into virtual velocity (rate), presenting the experience of flying. It is not generally considered natural because it requires completely different muscle groups from those used for physically moving through an environment by walking. Though practice may lead to automatic, or natural, action (Still & Dark, 2013), possibly due to neural changes in the motor cortex (Grafton, Schmitt, Horn, & Diedrichsen, 2008), such a rate-controlled interface does not provide consistent feedback on distance traveled.

More natural interfaces exist, including treadmills (Christensen, Hollerbach, Xu, & Meek, 2000; Darken, Cockayne, & Carmein, 1997; Iwata, 1999; Wang, Bauernfeind, & Sugar, 2003), the VirtuSphere (Medina, Fruland, & Weghorst, 2008), and redirected walking (Hodgson, Bachmann, & Waller, 2011; Razzaque, Kohn, & Whitton, 2001). These devices are position-control interfaces because a user's movements directly affect position in the VE. None of the interfaces is completely natural, because feedback from the pro-

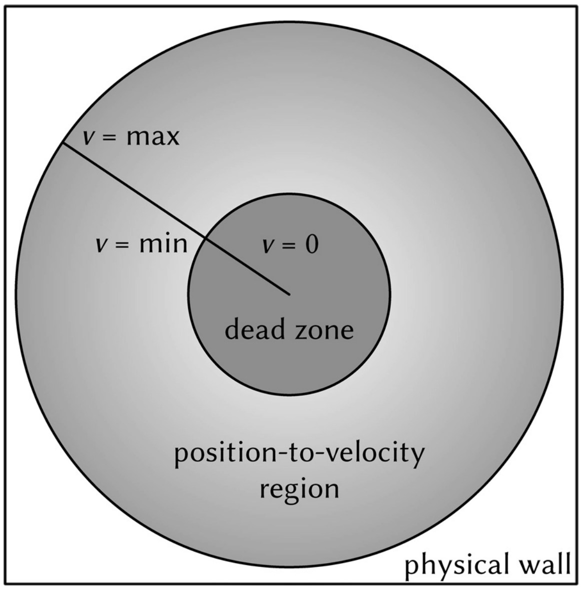

Figure 1. *Top-down depiction of the P2V locomotion interface. The user's velocity is indicated by* v.

prioceptive and vestibular senses is not the same as experienced while walking in the physical world. Also, some of these interfaces present problems for users, for example, when stopping or switching directions.

Hybrid interfaces implement position control to an extent, but as a tracked user approaches a system's boundaries, less natural rate-control techniques are used. One such interface is Magic Barrier Tape (Cirio, Marchal, Regia-Corte, & Lècuyer, 2009), which allows for completely natural walking near the center of the physical environment, but when the user nears the system boundaries, a virtual barrier tape must be crossed, at which point the interface becomes rate-controlled according to the user's position from the physical center.

A similar body-based interface, known as position-to-velocity, or P2V (depicted in Figure 1), was developed in the Virtual Reality Applications Center for use in the C6 CAVE, a six-side projection-based virtual-reality display. When a user of this interface is near the center of the CAVE (i.e., the dead zone), the interface is position-controlled, so the user can walk physically. Once the user steps outside of the dead zone, the interface becomes rate-controlled and the user's head position is used to calculate a virtual velocity vector with respect to the CAVE center. The user can increase or decrease speed by stepping farther from or closer to the CAVE center, and can change the direction of travel by stepping to the left or right. To stop, the user must return to the dead zone. This interface is particularly well suited for systems, such as the C6, that provide a 360° field of view, allowing for fully natural rotations, yet are limited to a small physical movement area.

1.2 Field of View

Vision is essential for the effective control of locomotion. A typical human has a 200° horizontal and 135° vertical field of view (FOV; Wandell, 1995). Many VR systems utilize head-mounted displays (HMDs) to display virtual scenes, due to their size, flexibility, and cost. However, most HMDs provide fairly limited FOV. For example, the popular NVIS nVisor SX HMD has an FOV of 47° × 38°. CAVE-like systems often have a FOV limited only by a user's visual abilities, although they are expensive, inflexible, and graphics resolution is still usually restricted. Wearing stereoscopic glasses limit this FOV to about 140° × 90°, but such a system still boasts a larger FOV than nearly all currently available HMDs.

Past research has shown that users in VEs do not interpret spatial information, such as distances, as accurately as humans in real-world scenarios (see Loomis & Knapp, 2003, for a review). However, it is unclear what aspects of VEs contribute to these phenomena (Thompson et al., 2004). A common result in such studies is that participants tend to underestimate distances in VR, possibly due in part to a restricted FOV (Kline & Witmer, 1996; Willemsen, Colton, Creem-Regehr, & Thompson, 2009), although some studies have shown no effect of FOV on distance perception (Creem-Regehr, Willemsen, Gooch, & Thompson, 2003; Knapp & Loomis, 2004; Messing & Durgin, 2005).

Because distance perception is important during virtual locomotion and navigation, FOV may be an important factor contributing to performance and for determining which cognitive strategies are employed. There is also evidence that peripheral vision is important during locomotion, and some studies have associated a limited FOV with navigation and memory performance deficits (Alfano & Michel, 1990; McCreary & Williges, 1998). However, other studies have found a limited effect of FOV on spatial updating (Pèruch, May, & Wartenberg, 1997).

When humans move through the world with unconstrained vision, they normally view the environment with multiple overlapping eye fixations. While moving about, integrating information between fixations might be unnecessary, as most information will still be available in the periphery (Dolezal, 1982). This implies that locomotion with a reduced FOV may require additional storage and integration of information, thereby adding cognitively demanding processes that are not normally required during locomotion. During unconstrained locomotion, humans use patterns of visual stimulation, known as optic flow, to extract information about movements and displacements relative to their environment (Gibson, 1986). The availability of this information has a profound impact on locomotion. For example, optic flow rate has been shown to impact the transition from walking to running (Mohler, Thompson, Creem-Regehr, Pick, & Warren, 2007). Humans are capable of locomoting without normal optic flow (Loomis, Beall, Macuga, Kelly, & Smith, 2006; Macuga, Loomis, Beall, & Kelly, 2006; Warren, Kay, Zosh, Duchon, & Sahuc, 2001), but this may lead to alternate strategies to judge distance traveled and orientation. Because restricting FOV leads to reduced optic flow, strategies employed during virtual locomotion may require additional cognitive resources.

1.3 Cognitive Resources

Working memory is thought to rely on a finite pool of cognitive resources. Most theoretical models of cognitive resources include multiple components of working memory and, at minimum, draw a distinction between verbal and spatial resources. Baddeley and Hitch (1974) proposed one such model, which includes a visuo-spatial sketchpad for manipulation and storage of visual and spatial information, and a phonological loop for manipulation and storage of phonological (i.e., linguistic speech-based) information. According to this model, a central executive controls general attention, which mediates access to the two working memory subsystems. The model has been expanded over the years (Baddeley, 2002), and some evidence now points to a further division of visuo-spatial resources into separate pools for visual and spatial tasks (Darling, Della Sala, & Logie, 2009).

Researchers use a dual-task selective-interference paradigm to assess the resource demands of a task of interest. In such studies, a participant performs an experimental task concurrently with a second task of known demands (in terms of a multicomponent model of working memory, as described above). If performance drops on either task, then it can be concluded that the experimental task of interest requires resources from the same pool as the known secondary task. In terms of the Baddeley and Hitch (1974) model, these resources would be either visuo-spatial or verbal. If both verbal and spatial tasks interfere similarly with the task of interest, then the task may require general attention resources or an equivalent amount of both verbal and spatial resources.

Past research has shown that using unnatural locomotion interfaces is cognitively demanding (Zanbaka, Lok, Babu, Ulinski, & Hodges, 2005). However, no previous study has attempted to examine the nature of these demands. In the studies that follow, the dual-task selective-interference paradigm is used to examine the impact of different cognitive tasks on specific aspects of locomotion. The first study compares the cognitive-resource requirements of interfaces varying in their naturalness. The second study examines the impact of a restricted FOV on cognitive-resource demands during locomotion.

2 Study 1: Working Memory Use During Seminatural Locomotion

Previous research has not examined the specific resource demands of unnatural locomotion, but it is reasonable to expect that spatial resources are required because locomotion is an inherently spatial task, and because unnatural interfaces require an accurate mental model of control movements that will lead to appropriate actions within the virtual environment. General attention resources may also be involved, possibly due to the novelty of locomotion, path-planning challenges, or reduced sensory feedback for error correction (Shallice, 1982). Additionally, it was expected that performance on the more unnatural aspects of an interface would suffer the most when placed in competition for finite resources. For example, the movements for stopping are unnatural in the P2V interface (requiring a return to the CAVE center), but rotating in place involves completely natural movement.

A study was conducted to investigate the impact of concurrent cognitive (working memory) tasks on movement performance using two locomotion interfaces: a variant of the body-based position-to-velocity technique described in Section 1.1 (P2V) and a standard gamepad interface (GP). As a baseline, the study also included a real-walking group (RW) in which participants physically moved about, just as in real life. This study aimed to determine if two types of cognitive tasks (spatial and verbal) had different impacts on movement performance when using three different locomotion interfaces along a spectrum of naturalness, in order from lowest to highest: GP, P2V, RW. The working memory tasks required participants to remember a sequence of either spatial or verbal items and then recall those items after completing locomotion tasks.

2.1 Methods

The following study design was approved by the Iowa State University Institutional Review Board.

2.1.1 Participants. Fifty-one undergraduate students were recruited from the Iowa State University Department of Psychology research participant pool (received course credit) and by word of mouth (received no compensation), and were randomly assigned to one of three experimental interface groups. Participants came from multiple departments and majors across campus. All participants were required to have 20/20 (corrected) binocular vision and all played less than or equal to 3.5 hr of first-person video games in a typical week. This requirement was included because participants with extensive gaming experience may be more familiar with interface technology. Additionally, participants in the GP interface group were further restricted to no more than one-and-a-half hours of first-person video games per week, in an attempt to limit gamepad familiarity. In this first study, the gamepad was intended to be representative of a typical unnatural locomotion interface, so it was important to ensure that it was, in fact, unnatural.

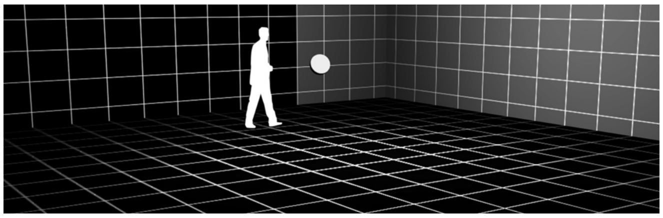

Figure 2. *Virtual room used in Study 2 and similar to that used in Study 1.*

2.1.2 Stimuli and Design. The experimental phase took place in a virtual grid room (displayed in the CAVE), similar to that rendered in Figure 2. The grid texture was intended to be simple, but to provide adequate visual feedback to guide locomotion. One wall of the virtual room was purple, and the participant was instructed to always face that wall with his or her body (head rotations away from the purple wall were allowed). As seen in the figure, translation tasks involved moving to the location of a virtual golden nugget, with a radius of 30.48 cm, centered 129.54 cm above the floor and 152.4 cm from the participant. For a given task, the nugget was displayed in front of or to either side of the participant. For the case where the nugget was displayed to the side, sidestepping was required in order to retrieve it (while continuing to face the purple wall). For rotation tasks, the room instantaneously rotated (i.e., there was no optic flow, just a discreet change) 90° to the left or right, and participants found themselves no longer facing the purple wall, necessitating a rotation to continue facing the purple wall. Finally, a virtual I-beam occasionally flew toward the participant, 152.4 cm above the floor, requiring ducking to avoid being hit.

VirtuTrace, a generalized VR experiment engine designed for flexible configuration and swapping of interfaces, controlled all stimuli and handled interface inputs. Each participant was randomly assigned to one of the following three conditions, listed in order from least to greatest expected naturalness:

Gamepad (GP): The participant used a Logitech Wingman gamepad. The left stick was used for virtual translation while the right stick was used for rotation. Full-body physical rotation was not allowed, so the participant stood in the center, facing the front wall of the CAVE for the entire study. The A button controlled ducking, but the participant was also allowed to duck physically, as detected by head position, if preferred.

Body-based Position-to-Velocity (P2V): The participant's physical head position was used to set a velocity in the virtual world. Instead of a single velocity vector originating at the center of the CAVE, as described in the introduction, there were two separate velocity vectors: one representing the distance from each axis. Because all of the required movements were axis-aligned, this did not change the optimal movements. Calculating a participant's velocity in this way reduces the chance of drifting off course on axis-aligned tasks, and could therefore make the movements easier. The dead zone for each axis was set to 42.67 cm in CAVE space. All movements were 1 : 1 (i.e., a translation of 30 cm in the physical world corresponded to an identical translation of 30 cm through the virtual world) within the dead zone. Rotation and ducking were performed physically, as in everyday movement.

Real-Walking (RW): The participant moved around the CAVE using real-life walking, rotation, and ducking. All movements were 1:1.

No single interface was expected to be "best" in all cases because each may be useful for different types of tasks under different concurrent task conditions. The following tasks were performed in each round: translate left, translate right, translate forward, rotate left, rotate right, and duck.

Translate Left: The participant retrieved a golden nugget to the left. An arrow appeared on the floor in front of the participant, indicating the location of the nugget. Because the purple wall was still in front of the participant, sidestepping was required to correctly complete the task.

Translate Right: The participant retrieved a golden nugget to the right. An arrow appeared on the floor in front of the participant, indicating the location of the nugget. Because the purple wall was still in front of the participant, sidestepping was required to correctly complete the task.

Translate Forward: The participant retrieved a golden nugget to the front.

Rotate Left: The environment rotated such that the purple wall was on the left side of the participant, cuing a 90° rotation to the left in order to continue facing the purple wall.

Rotate Right: The environment rotated such that the purple wall was on the right side of the participant, cuing a 90° rotation to the right in order to continue facing the purple wall.

Duck: The participant had to duck to avoid being hit by a virtual I-beam.

Note that the participant was instructed to stop immediately after completing the translation tasks. In the RW group, this required the participant to stop in place, then return to the CAVE center in preparation for the next task.

If unnatural locomotion requires additional cognitive resources, the participant should exhibit lower performance on locomotion actions and/or working memory tasks when both are presented concurrently. If spatial resources are used more than verbal resources, then the results should show a decrease in performance associated with a simultaneous spatial task. Conversely, if verbal resources are required, then the results should show a decrease in performance associated with a simultaneous verbal task. If general attention resources are required, then a similar performance decrease should be seen with either a verbal or a spatial concurrent task.

Performance was examined independently for specific aspects of locomotion: translating forward, sidestepping, turning, and ducking. The naturalness of an interface depends on the action performed, and performance on less-natural actions should be more affected by the addition of a concurrent cognitive task. For example, the P2V interface allows for completely natural rotation within the dead zone, so there should be no detriment in rotation performance when a concurrent cognitive task is added. However, stopping with that same interface is unnatural, requiring a participant to locate and return to the center of the CAVE. This should lead to decreased performance in terms of stop time and possibly in the ability to remember the memory items (if competition does exist).

VirtuTrace tracked and logged the participant's head position, measured by an InterSense IS-900 tracking system. During data analyses, a moving average of head positions was used to automatically determine when the participant had stopped (therefore signaling the end of the locomotion task). The following performance metrics were recorded or calculated for each task:

Locomotion Time: The time from locomotion task presentation until completion of each translation or rotation task.

Start Time: The time from task presentation until movement started for each translation task.

Stop time: The time from task completion until movement stopped after each translation task.

Path Length: The length of the traversed path for each translation task.

Duck Failure/Success: Ducking was successful if the participant was not hit by the beam.

Memory Items Missed: The number of items missed in the working memory task recall phase, calculated as the minimum number of swaps or replacements needed to transform the response into the correct response.

2.1.3 Procedure. Upon arrival, the participant completed a questionnaire involving demographic information, video game experience, and participation in athletic activities. Next, the Perspective Taking/Spatial Orientation Test (PTSOT; Hegarty & Waller, 2004) was administered. In the PTSOT, the participant is given 5 min to make a series of 12 egocentric direction judgments based on imaginary orientations in simple scenes drawn from an overhead perspective (e.g., "Imagine you are standing at the stop sign, facing the cat. Draw a line in the direction of the house."). Because effective locomotion requires users to make predictable changes to their egocentric perspectives, the PTSOT may be a useful way to examine individual differences in the resource demands of locomotion.

The participant then entered the C6, a fully immersive CAVE with four walls, floor, and ceiling allowing a 3.05 m^2 movement area, illuminated from outside with stereoscopic computer graphics. The participant was given instructions and a demonstration of how to complete working memory tasks in the VE. Still in the CAVE, he or she then completed a series of six verbal working memory trials with immediate recall. These trials were of increasing difficulty and were intended to roughly assess individual verbal spans and allow practice on the verbal working memory task. The difficulty increased incrementally from three to five items, with two trials at each difficulty level. Next, the participant was trained on the spatial working memory task and given a series of six spatial trials, again increasing in difficulty from three to five items. Trials at the lowest span (three) were intended to provide practice, and trials at the highest span (five) were used to customize the difficulty of the verbal and spatial working memory tasks during the experimental blocks. When a participant was unable to successfully complete a verbal or spatial task at the highest difficulty level (five), then the span used for that particular task type (verbal or spatial) was dropped to four when performed concurrently with the locomotion task. This was done to ensure that the span used during the locomotion tasks was sufficiently large to tax the cognitive resource in question, but not so large that the participant was incapable of recalling the span. The preassessment trials also provided practice so that the participant would feel comfortable with the memory tasks in the experimental blocks.

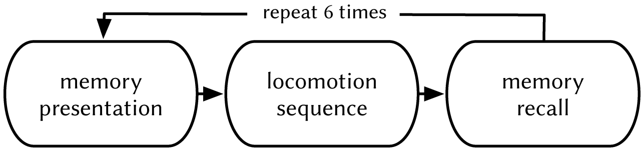

Figure 3. *Task flow of a single experiment block.*

Before the experimental blocks, the participant was given instructions and a detailed demonstration of how to complete all locomotion tasks using the assigned interface. However, the participant only observed and was not allowed to use the interface during the demonstration. This served to maintain the unnaturalness of the interface and maximize the extent to which cognitive resources would be required when using it. All tasks were performed in the grid room described in Section 2.1.2. The participant was instructed to always face the purple wall and to stand still in the center of the CAVE between tasks.

The study was composed of six blocks of locomotion tasks with concurrent working memory tasks, as diagrammed in Figure 3. In each block, a participant was presented with a working memory span sequence, followed by a series of locomotion tasks, and finally asked to recite the working memory sequence. Therefore, the participant had to maintain the verbal or spatial span in working memory while performing the locomotion tasks. Each block had a verbal, a spatial, or no working memory task assigned randomly, such that each participant performed two blocks under each working memory condition. Each block lasted at least 70 s, which was more time than was typically required to complete the locomotion tasks. This was done to reduce incentive for the participant to rush through the locomotion task in order to reach the recall phase more quickly. The sequence of locomotion tasks was randomly ordered within each block. Each of the locomotion tasks listed in Section 2.1.2 was performed once during each experimental block.

Between locomotion tasks, there was a 6-s pause, during which the participant stood in the center of the CAVE and awaited the next locomotion task. The study contained six blocks of the events described above, two blocks with each working memory task.

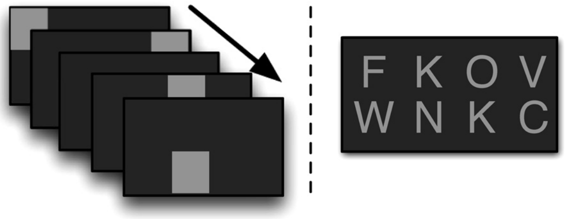

Figure 4. *Sample spatial memory task presentation sequence (left) and recall card (right).*

After completing all six blocks of trials, the participant was asked to complete a post-questionnaire and answer questions in an unstructured interview. These questions were intended to uncover any strategies used or problems encountered, specifically involving the participant's perceived competition for cognitive resources.

For the verbal and spatial memory tasks, a memory sequence was presented before the participant performed the locomotion movements. The participant was later asked to recall the sequence after the locomotion movements were complete. For the verbal task, the presentation phase involved a sequence of numbers, and during the recall phase, the word "recite" appeared, instructing the participant to recall the numbers. A simple one-word recite card was designed so that it would not interfere with a participant's ability to remember the verbal sequence. The flow of spatial memory presentation and recall is shown in Figure 4. For spatial tasks, a matrix of random letters was displayed and the participant was required to state the letters that corresponded to the positions, in order, in which the boxes were previously displayed. This recall method was devised because participants in the GP group were already using a gamepad for locomotion actions; therefore, using a gamepad or other hand-operated input device for the memory recall might have led to confusion or interference between the tasks. The recall letters were randomized to prevent the participant from encoding the spatial locations verbally, which would have taxed the wrong working memory resource. For the control task (none), the participant was instructed to stand still in the center of the CAVE while the word "Wait" was displayed during both presentation and recall.

2.2 Results

The following analyses focus on the effects of the locomotion interfaces and memory tasks on movement performance (locomotion time, start time, stop time, and duck failure rate) and on working memory performance (memory items missed). On many translation trials, participants took suboptimal paths, which led to the movements not being atomic. For example, if a participant sidestepped left but missed the nugget, a right sidestep (or some other combination of movements) was required in order to eventually reach the nugget and complete the task. Therefore, data from left, right, and forward translation trials were combined prior to analysis. It is not appropriate to interpret the previously described movement as a left-sidestep action as initially intended. Rather, it is more reasonable to treat it as a single translation movement, without regard for direction.

Stop time, duck failure, and working memory items missed were most responsive to manipulations of interface and/or memory task, so the following analyses focus on those measures. Additionally, a performance decrement on either the movement tasks or concurrent memory tasks is considered evidence of competition for working memory resources, meaning that the tasks had overlapping resource demands. A condensed version of the analyses of stop time and working memory performance was described by Marsh, Putnam, Kelly, Dark, and Oliver (2012). This section expands on those results, in particular adding detail regarding start time findings and including an analysis of duck failures.

For some trials, recorded data points were removed or did not exist for one of the following reasons:

- In many locomotion trials, the participant was not fully stopped before the next task was presented; therefore, stop time (for the previous trial) and start time (for the subsequent trial) were not recorded.
- Hardware and software problems led to incomplete data for some participants.
- In a few cases, the participant missed the nugget but thought it had been retrieved. Because the objective was to measure the ability to successfully complete intended movements, head position data were manually inspected and discarded where it was clear that the participant had passed the nugget and stopped, preparing for the next task, before realizing the mistake.
- Some participants reported using a verbal strategy to remember spatial sequences (i.e., assigning numbers to spatial positions). Because the task was intended to tax spatial resources, affected data were removed whenever a participant reported using such a strategy.

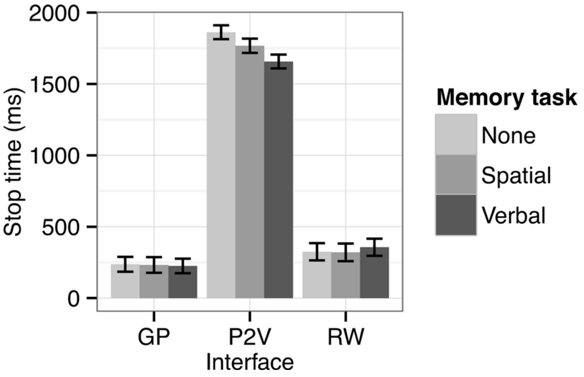

Figure 5. *Study 1 mean stop time as a function of interface (GP = gamepad, P2V = position-to-velocity, RW = real-walking) and memory task. Error bars show ±1 standard error of the mean.*

Across all Study 1 analyses, the percentage of data missing or removed for these aforementioned reasons ranged from 7.3% to 24.7%, with an average of 14.7%. There was no indication of patterns among these data points.

2.2.1 Stop Time.

Stopping is an aspect that differs greatly between interfaces. In the RW group, participants must simply stop walking. Participants in the GP group can stop by releasing the stick on the gamepad. To stop with the P2V interface, participants were required to return to the center of the CAVE. The latter should be more time-consuming and it may require cognitive resources for orientation and necessary movements.

Average stop times for all translation tasks and interfaces are shown in Figure 5. A mixed-model ANOVA was performed with fixed effects for interface group

(GP, P2V, RW) and working memory task (spatial, verbal, none), and random effects for participants. Significant main effects of locomotion interface, $F(2, 45) = 298.74, p < .001$, and memory task, $F(2, 641) = 4.22, p = .015$, were qualified by a significant interaction between locomotion interface and memory task, $F(4, 641) = 3.36, p = .01$. The main effect of interface groups was expected, because stopping with the gamepad (let go of the stick) or real-walking (stand still) interface is trivial, while stopping with the P2V interface requires locating and returning to the CAVE center. This expectation was supported by the analysis. Also, because stop times were so low in the GP and RW groups, one should not expect to see a difference between memory tasks within those groups. This was also supported by the analysis. A Markov chain Monte Carlo (MCMC) simulation from the posterior distribution for the model was used to obtain estimates and p values for comparisons of interest. The most interesting stop time results are found in the P2V group. Participants using this interface stopped significantly faster when performing a spatial working memory task than when performing no task ($p = .04$), and significantly faster when performing a verbal working memory task as opposed to a spatial working memory task ($p = .02$).

The faster stop time when performing a spatial or verbal working memory task compared to no working memory task may initially seem counterintuitive. However, one explanation may be that participants were motivated to move faster in order to end the resource competition between the locomotion task and the working memory task. Although a minimum of 70 s elapsed between presentation and recall of the working memory span, it was the movements that were expected to compete with maintenance of the memory items, not the elapsed time between span presentation and recall. Therefore, the participants may have minimized locomotion times in order to reduce resource competition. This conclusion is supported by participant feedback, which indicated a subjective sense that the locomotion and working memory tasks competed for resources. It is also supported by the start time results described below.

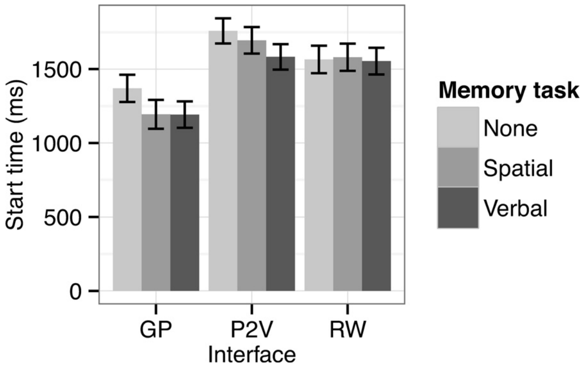

Figure 6. *Study 1 mean start time as a function of interface (GP = gamepad, P2V = position-to-velocity, RW = real-walking) and memory task. Error bars show ±1 standard error of the mean.*

The difference in stop time when performing spatial and verbal tasks is an intriguing result. There are at least two possible explanations for this difference. First, participants might have been more motivated to complete the locomotion task quickly when performing a concurrent verbal task compared to a spatial task, presumably because they felt increased competition for verbal resources. However, it is unclear why verbal resources would be taxed more than spatial resources, so this explanation does not fit well with the theoretical premise of the study. Second, participants might have been equally motivated to complete the locomotion task quickly during both types of memory task, but they may have been incapable of stopping as fast during the spatial task due to competition for spatial resources. The second explanation fits traditional dual-task interpretations, and is supported by a visual inspection of the start times depicted in Figure 6, although there is no significant interaction in the start-time data. Further evidence supporting the second stop-time interpretation is seen in the start-time data of the GP group, where performance with no memory task is slower and there is no meaningful performance difference during a concurrent spatial or verbal task. Because stopping is trivial when using a gamepad (one must simply let go of the stick), the stop time data do not show this trend. These patterns support the notion that participants were equally motivated when performing spatial and verbal tasks. Self-reported feedback from post-questionnaires and exit interviews also supports the interpretation that spatial memory tasks interfered with

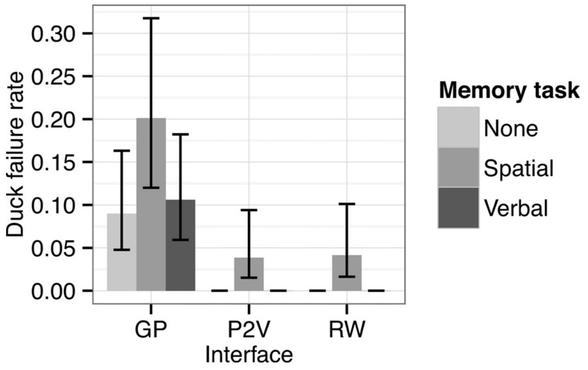

Figure 7. *Study 1 mean duck failure rate as a function of interface (GP = gamepad, P2V = position-to-velocity, RW = real-walking) and memory task. Error bars show ±1 standard error of the mean.*

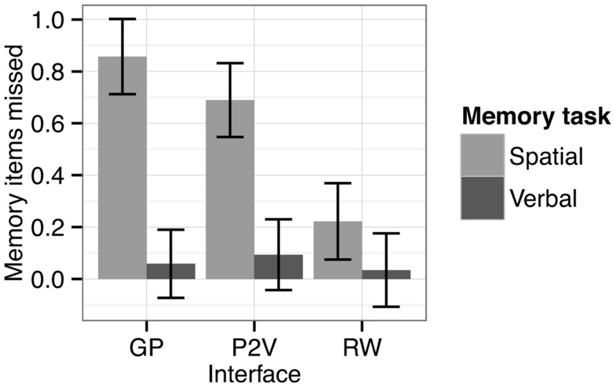

Figure 8. *Study 1 mean number of memory items missed as a function of interface (GP = gamepad, P2V = position-to-velocity, RW = real-walking) and memory task. Error bars show ±1 standard error of the mean.*

movement performance to a greater degree than did verbal memory tasks.

2.2.2 Duck Failure Rate. Failure rates on the ducking task are plotted in Figure 7. Recall that the ducking task required the participant to duck to avoid an overhead I-beam. A failure on this task is defined as being hit by the beam. Because there were zero failures for some combinations of independent variables, two single-factor mixed-model analyses were performed, treating failures as binomially distributed. These analyses showed a significant main effect of interface, $F(2, 45) = 4.87, p = .012$, and a marginally significant main effect of memory task condition, $F(2, 214) = 2.75, p = .066$. The failure-rate data show that participants had particular problems using the gamepad for ducking. Recall that participants using the gamepad interface were allowed to either duck physically, as detected by head position, or by pressing a button on the gamepad. Because GP participants frequently used both methods simultaneously, no reliable data are available regarding preference for one over the other. During the study, experimenter observations revealed that participants commonly made the mistake of releasing the button prematurely, causing them to stand up before the beam had completely passed. Additionally, although the main effect of memory task did not reach significance, duck failure rates were nominally higher when performing a concurrent spatial task as compared to a verbal task or no task, even when performing the action as they would in real life.

2.2.3 Memory Items Missed. Multicomponent models of working memory (e.g., Baddeley & Hitch, 1974) predict a performance detriment on either task, or both, when simultaneously performing two tasks requiring common cognitive resources. The number of missed items on the memory tasks is plotted in Figure 8. A two-factor ANOVA with terms for interface (GP, P2V, RW) and working memory task (spatial, verbal) revealed a significant main effect of memory task, $F(1, 40) = 20.609, p < .001$, a marginally significant main effect of interface, $F(2, 40) = 2.933, p = .065$, and a marginally significant interaction between memory task and interface, $F(2, 40) = 2.494, p = .095$. A possible interpretation of the significant main effect of memory task is that the spatial tasks were simply more difficult than the verbal tasks and so more items were missed. However, all participants remembered 100% of items on both types of memory tasks during the preassessment, which included no concurrent locomotion task, up to their respective individual spans. Perhaps the 70-s retention time during the study selectively increased the difficulty of the spatial task. However, based on the expected results, the overall patterns in the data, and the self-reported participant feedback, the working memory performance difference is likely due to resource competition from the concurrent locomotion tasks. For this reason, spatial working memory task performance was analyzed in isolation. It was expected that the gamepad would be the least natural interface, the position-to-velocity interface would be more natural, and real-walking would be completely natural. Based on this a priori hypothesis, the pattern across the GP, P2V, and RW interfaces was analyzed using predicted contrast weights $(1, 0, -1)$. The contrast was found to significantly describe the data, $F(1,41) = 5.394$, $p = .03$.

2.3 Discussion

The findings above indicate that locomotion competes for spatial working memory resources, with some locomotion interfaces imposing higher demands. Specifically, a simultaneous spatial memory task hindered stopping performance when using the position-to-velocity interface to perform locomotion tasks, which required finding and returning to the center of the physical environment. Conversely, participants exhibited decreased spatial working memory performance while concurrently performing unnatural locomotion actions. Such problems were not present for other tasks, such as rotation, that are more natural. Finally, although the results did not reach statistical significance, participants also exhibited problems with ducking to avoid obstacles while performing a concurrent spatial memory task, even when using their full bodies to duck, just as they would in an everyday physical environment. This possibility warrants further research.

The metrics that did not respond to experimental manipulations may be those interface aspects that are more natural and/or easy. For example, it should not be difficult to move in a straight path with any of the interfaces, so we can expect a lack of difference in path length. Additionally, in this study, participants appeared to make the trade-off in favor of locomotion performance over memory performance. If a study were designed to place greater importance on the cognitive task, performance would perhaps suffer on locomotion metrics.

Users of VR systems are often asked to perform multiple simultaneous tasks. Many of these tasks impose extreme cognitive demands, and they can be critical to success in the underlying scenario. Examples of such dual-task applications include firefighter training and military command and control scenarios. In such scenarios, problems associated with concurrent tasks may be compounded with the addition of stress and fatigue. The findings from the current study, together with domain-specific knowledge of common cognitive tasks, could be used to inform the design of future VR systems, particularly the choice of locomotion interfaces.

Together, unnatural control actions coupled with reduced sensory feedback compose an unnatural locomotion interface. This study investigated cognitive demands related to unnatural actions required to effect locomotion in VR. The next study will address the problem of limited sensory feedback, specifically reduced FOV, provided by such an interface during locomotion.

3 Study 2: Working Memory Use with a Restricted Field of View

A second study was conducted to investigate the impact of a restricted FOV on cognitive resource demands while using a virtual locomotion interface. The reduced visual feedback associated with a reduced FOV was expected to cause participants to resort to alternative locomotion strategies requiring additional working memory resources. Spatial resources might be required if the alternative strategies depend on mentally storing and manipulating spatial information that is normally visible in the environment. Verbal resources might be required if the alternative strategies depend on verbal coding of information such as distance traveled (e.g., counting steps). General attention resources may also be involved when FOV is reduced, possibly for planning or error correction. If spatial resources are required, locomotion and/or memory task performance should decrease when given a concurrent spatial task and a reduced FOV. If verbal resources are required, locomotion and/or memory task performance should decrease when given a concurrent verbal task and a reduced FOV. If general attention resources are required, then an equal decrease should be exhibited in locomotion or memory task performance with either a verbal or spatial concurrent task.

Because the second study was intended only to examine the impact of FOV, all participants used the same interface. The P2V interface was chosen because it is more natural than the gamepad, yet it allows for virtual movement beyond the CAVE boundaries.

3.1 Methods

The following study design was approved by the Iowa State University Institutional Review Board.

3.1.1 Participants. Thirty-one undergraduate students (20 males) were recruited from the Iowa State University Department of Psychology research participant pool (received course credit) and by word of mouth (received no compensation). All participants were required to have (corrected) 20/20 visual acuity. They were assigned randomly to two FOV groups. The high-FOV group wore CrystalEyes stereo shutter glasses, providing a FOV of approximately 140° × 90°. The low-FOV group wore the same glasses, but with cardboard attached in front of the lenses, reducing the FOV to approximately 60° × 45°.

3.1.2 Stimuli and Design. The study incorporated a 2 × 3 design, with FOV (high, low) as a between-participants variable and memory task type (spatial, verbal, none) as a within-participants variable. All participants used the P2V interface described above in Section 1.1.

All locomotion tasks took place in a virtual room (depicted in Figure 2) similar to that used in Study 1. Participants performed three basic types of locomotion tasks: translation, rotation, and ducking. Because participants in Study 1 generally completed the tasks quickly, two translation tasks of each type (left, right, forward) were included in each block in Study 2. As in Study 1, one of each rotation task (left, right) and one ducking task was included in each block in Study 2. Tasks occurred in a random order within each block. As in Study 1, a memory span task (spatial, verbal, none) was presented at the beginning of each block and recalled at the end of each block, such that there were two blocks of each type of memory task, randomly ordered.

Figure 9. *Staggered positions for spatial memory items in Study 2.*

Because there was no real-walking interface group in Study 2, there was no need for the nugget to be reachable within the bounds of the CAVE. For this reason, the nugget used in translation tasks was centered 213.36 cm from the participant, completely outside the physical walls of the CAVE. Also, to increase ducking task difficulty compared to Study 1, the I-beam height was lowered to 144.78 cm.

To deter participants from attempting to encode spatial memory items verbally, the memory task presentation cards were altered from Study 1 such that the items were staggered spatially as seen in Figure 9.

The same variables were recorded as in Study 1. Only significant findings are reported below. Because Study 2 was not intended to replicate Study 1, the analyses were not identical.

3.1.3 Procedure. The participant first completed a questionnaire including demographic information and questions about video game experience. Then the PTSOT was administered to investigate the role of individual differences in spatial ability on the cognitive demands of a limited FOV.

Next, the participant entered the CAVE and was trained on how to complete verbal memory tasks. Then the participant completed a series of verbal practice tasks. The practice tasks increased in difficulty from three to six items, with two trials at each difficulty level. As in Study 1, a participant's performance on the practice tasks was used to customize the task difficulty during the experimental blocks. Note that the highest possible memory span was increased from Study 1, to six items, in order to provide greater customization to a participant's abilities, because most participants in Study 1 had a span of five items.

The participant was then trained on how to complete the spatial memory tasks. This was followed by a series of spatial practice tasks, administered exactly as the verbal practice tasks.

Next the experimenter gave a demonstration of how to complete the locomotion tasks in the VE. As in Study 1, the participant was not allowed to practice using the interface, but simply observed the demonstration. The six experimental blocks then proceeded as in Study 1.

3.2 Results

As in Study 1, hardware problems, software problems, participant confusion, and self-reported verbal recoding of spatial locations led to data points being discarded. Across all Study 2 analyses, the percentage of data points missing or removed ranged from 2.9% to 23.5%, with an average of 10.6%. Additionally, one participant's data were removed completely because he reported disobeying directions and not taking the tasks seriously. There was no indication of patterns among these removed data.

A condensed version of these results was described by Marsh, Kelly, Dark, and Oliver (2012). This section expands on those analyses and adds consideration of stop time.

3.2.1 Stop Time. Stop time did not vary as a function of condition in Study 2, but each mean was faster than its counterpart in the P2V group of Study 1. One possible reason that stop times did not respond to experimental manipulations in Study 2 but did respond to similar manipulations in Study 1 is that participants had more time to plan their stop due to the longer distance traveled in Study 2 compared to Study 1. This could also have led to the overall faster stop times in Study 2 compared to Study 1.

3.2.2 Start Time. Start time reflects the time required for identifying the task to be performed and for planning and initiating the movement. A reduced FOV leads to reduced visual feedback during movement. When a participant expects reduced feedback, strategies may shift to involve more planning. Because FOV affects optic flow differently for left/right translations as compared to forward translations, analysis was done separately for left/right (sidestepping) translation tasks and forward translation tasks. A plot of the left and right (sidestepping) start time means, shown in Figure 10, provides some support for the conclusion that virtual locomotion with a restricted FOV requires additional general attention resources, although start time was not significantly influenced by the experimental manipulations in Study 2. Verbal and spatial times both showed a slight but nonsignificant increase when the FOV was reduced.

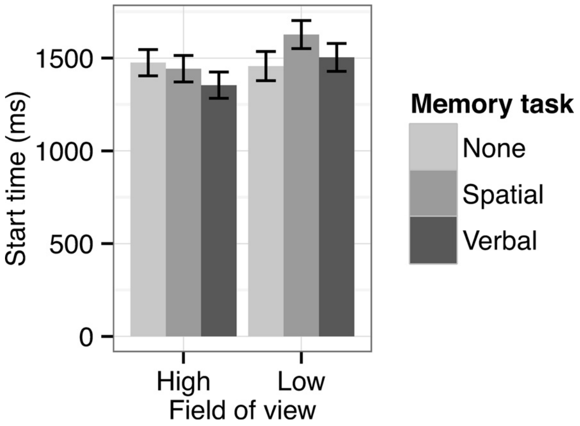

Figure 10. *Study 2 mean start time for left and right translation tasks as a function of field of view and memory task. Error bars show ±1 standard error of the mean.*

3.2.3 Memory Items Missed. Because there were more memory items in Study 2 than in Study 1, the number of items missed on the working memory task was treated as a Poisson distribution. A two (FOV: low or high) by two (memory task: verbal or spatial) mixed-model analysis, with random effects for subjects, on those data showed significant main effects of FOV, $F(1,27) = 4.27, p = .049$, and memory task, $F(1,69) = 25.26, p < .001$. The means are plotted in Figure 11. Restricting the FOV led to a similar

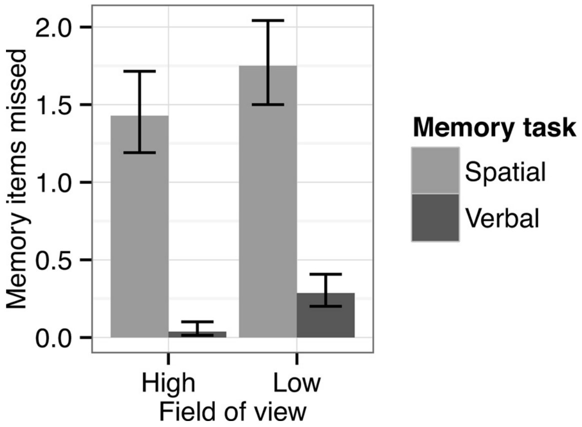

Figure 11. *Study 2 mean number of memory items missed as a function of field of view and memory task. Error bars show ±1 standard error of the mean.*

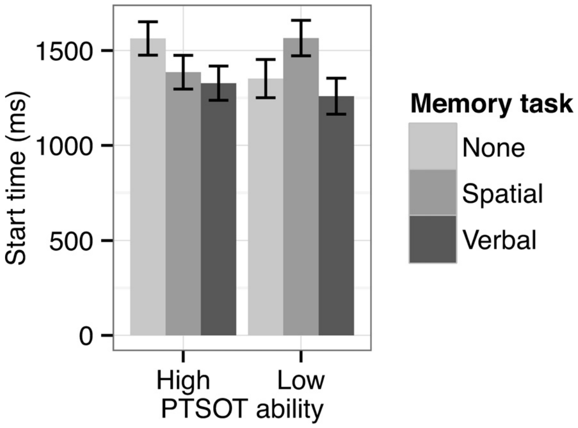

Figure 12. *Study 2 mean start time for left and right translation tasks as a function of PTSOT ability and memory task. Error bars show ±1 standard error of the mean.*

increase in memory items missed for both verbal and spatial memory tasks. In terms of the multicomponent models of working memory, this pattern of results indicates that additional general attention resources were required when the FOV was reduced. Alternately, an equal amount of verbal and spatial resources may have been required, which would also explain the observed results. Under either interpretation, there was no evidence that FOV reduction selectively required additional verbal or spatial resources.

3.2.4 Perspective Taking/Spatial Orientation Test. Answers on the PTSOT were scored by averaging each participant's deviation from the correct answer on all attempted items. Only 4.89% of the questions were not attempted in this study. As described by Kozhevnikov, Motes, and Hegarty (2007), participants with average errors in the bottom quartile ($8.58°-14.5°$) were placed in the high ability category (7 males, 2 females) and those with average errors in the upper quartile ($37.82°-109.11°$) were placed in the low ability category (5 males, 3 females). All participants in the second and third quartiles were eliminated from the following analyses.

Participants with low perspective-taking ability started sidestepping tasks more slowly when concurrently remembering a spatial sequence than when remembering a verbal sequence or no sequence at all. A two (FOV: low or high) by two (memory task: verbal or spatial) by two (PTSOT: high or low) mixed-model ANOVA, with random effects for subjects, was conducted on the left and right translation start time. A significant main effect of memory task, $F(2, 263) = 5.72, p < .004$, was qualified by a significant interaction between memory task and PTSOT ability, $F(2, 263) = 4.44, p = .01$. Figure 12 shows a plot of the means. This makes sense, as an individual with a lower spatial ability should be expected to perform worse on spatial tasks, and it is reasonable to expect that planning and initiating bodily movements requires spatial resources. In contrast, participants with high perspective-taking ability did not exhibit the same detriment from a concurrent spatial task. The primary difference across memory task conditions in the high spatial ability group is the slower performance when given no concurrent memory task. This may mean that users with high spatial abilities use more time-consuming strategies when planning and initiating locomotion movements, if resources are not already being used for another task. Alternately, note that this pattern closely resembles that seen in the Study 1 start and stop times, described in Section 2.2.1. Accordingly, participants with high spatial abilities may be more motivated to finish the task to reduce interference.

3.3 Discussion

The results presented in Study 2 provide evidence that a reduction in FOV leads users to resort to locomotion strategies that require additional general attention resources. In the Baddeley (2002) multicomponent model of working memory, this would implicate the central executive. General attention resources of the central executive may be involved when FOV is reduced, possibly due to path-planning challenges associated with a reduced FOV or reduction in error-corrective visual feedback. Alternately, these results do not eliminate the possibility that a similar amount of verbal and spatial resources are required for locomotion with a reduced FOV.

The PTSOT results indicate that individuals with different spatial ability levels are differentially impacted by the addition of a concurrent cognitive task. Specifically, individuals with low perspective-taking ability showed decreased performance when performing locomotion tasks with a concurrent spatial memory task, and individuals with high perspective-taking ability took longer to plan and initiate sidestepping movements when there was no competition for resources. This latter result may indicate that users with high spatial ability spend more time planning movements when the required resources are available, or it may reflect a difference in motivation to end competition for resources.

These findings should inform the design of VR systems. When selecting a display technology, it is important to consider the types of locomotion activities to be performed, the presence of concurrent cognitive tasks, and also the spatial abilities of the user. The fact that participants in this study sacrificed performance on the memory tasks, as opposed to the movement tasks, reinforces the importance of these decisions, as the study tasks were intended to simulate the existence of real-world primary tasks. However, users may make sacrifices on locomotion tasks instead of the primary (i.e., cognitive) task if the primary task were of real-world importance. Future work should examine this trade-off. As in Study 1, changes that affect participant priorities may cause additional locomotion metrics to gain significance.

Future work should specifically investigate a possible effect of concurrent task load on translation start times, as these times may reflect differences in time spent planning. Such research should also attempt to determine what role individual differences play when such competition is present. Future studies can incorporate additional measures of individual abilities, as well as translation tasks designed to reveal more information on planning strategies in use.

4 General Discussion

Past research has shown that there are cognitive costs when using unnatural interfaces (Zanbaka et al., 2005). However, that research has not investigated the nature of those costs or what types of concurrent cognitive tasks might be impacted when using such interfaces. The current project was designed to address what can be thought of as the input and output limitations of a virtual interface, with the input being the control layout itself and the output being sensory feedback.

Study 1 findings indicate that additional spatial resources are required when manipulating an unnatural locomotion interface, such as a gamepad, compared to full physical movement, as required in the real-walking interface. In this context, seminatural interfaces, such as the P2V interface, provide a compromise between cognitive costs and movement restrictions. Past research has not considered the cognitive costs of individual interface actions. This information can be useful when designing a VR system for use in a domain that is highly spatial in nature. Alternately, if an interface must be chosen for a domain that is primarily verbal, then the choice of interface may be less critical. Further, this interference between spatial tasks and unnatural interfaces hinders completion of the user's primary task as well as the locomotion movements.

The findings from Study 2 indicate that reduced visual feedback due to reduced FOV creates additional demands for general attention resources. This conclusion means that FOV during locomotion is important for successful completion of concurrent tasks that require spatial, verbal, or general resources. Further research is required to examine the cognitive demands of reduced feedback for other senses, such as proprioceptive, vestibular, or auditory, as well as other aspects of visual feedback including stereo, resolution, and realism. Additionally, the findings from Study 2 show that the cognitive demands of locomotion are impacted by individual differences, possibly due to differing strategies. Future work should consider this.

It is likely that practice using an interface will increase the degree to which it is natural, and thus decrease the cognitive demands. The studies described here involved novice users and the single-session design did not allow for statistical measurement of a training effect. Future work could include a longitudinal study to investigate whether dual-task performance using a novel interface eventually equals that of the real-walking interface.

The dual-task selective-interference paradigm has proven useful in addressing these research questions, and it may be an appropriate choice for some of the future work described above. The method for presenting and recalling spatial working memory items worked well without the use of hands. Future researchers should be aware of participant motivation issues that were encountered in these studies and try to avoid them or incorporate them when drawing conclusions.

5 Conclusion

These studies showed specific cognitive demands of locomotion using common VR interfaces. Both the unnatural interface actions and the limited FOV provided by VR systems were shown to have unique cognitive impacts, providing insight into the cognitive strategies employed and predicting performance problems on specific types of concurrent tasks performed in virtual worlds. The results show that spatial resources were required for unnatural control movements and general attention is likely to be required for strategies resulting from a restricted FOV. These findings, together with domain-specific knowledge, should be used to inform the selection of locomotion interfaces and display technology.

Acknowledgments

This research was supported by grants from the US Air Force Research Laboratory (Grant Numbers FA9550-05-1-0384 and FA9550-07-1-0475). The authors gratefully acknowledge this support.

References

Alfano, P. L., & Michel, G. F. (1990). Restricting the field of view: Perceptual and performance effects. *Perceptual and Motor Skills*, *70*(1), 35–45.

Baddeley, A. D. (2002). Is working memory still working? *European Psychologist*, *7*(2), 85–97. doi: 10.1027//1016-9040.7.2.85

Baddeley, A. D., & Hitch, G. (1974). Working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 8, pp. 47–89). New York: Academic Press.

Christensen, R. R., Hollerbach, J. M., Xu, Y., & Meek, S. G. (2000). Inertial-force feedback for the treadport locomotion interface. *Presence: Teleoperators and Virtual Environments*, *9*(1), 1–14. doi: 10.1162/105474600566574

Cirio, G., Marchal, M., Regia-Corte, T., & Lécuyer, A. (2009). The magic barrier tape: A novel metaphor for infinite navigation in virtual worlds with a restricted walking workspace. *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology, VRST 2009* (pp. 155–162). New York: ACM. doi: 10.1145/1643928.1643965

Creem-Regehr, S. H., Willemsen, P., Gooch, A. A., & Thompson, W. B. (2003). The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual environments. *Perception*, *34*(2), 191–204.

Darken, R. P., Cockayne, W. R., & Carmein, D. (1997). The omni-directional treadmill: A locomotion device for virtual worlds. *Proceedings of 10th Annual Symposium on User Interface Software and Technology, UIST 1997* (pp. 213–221). New York: ACM. doi: 101145/263407.2635550

Darling, S., Della Sala, S., & Logie, R. H. (2009). Dissociation between appearance and location within visuo-spatial working memory. *The Quarterly Journal of Experimental Psychology*, *62*(3), 417–425. doi: 10.1080/17470210802321984

Dolezal, H. (1982). *Living in a world transformed: Perceptual and performatory adaptation to visual distortion*. New York: Academic Press.

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J. J. (1986). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Grafton, S. T., Schmitt, P., Horn, J. V., & Diedrichsen, J. (2008). Neural substrates of visuomotor learning based on improved feedback control and prediction. *NeuroImage*, *39*(3), 1383–1395.

Hegarty, M., & Waller, D. (2004). A dissociation between mental rotation and perspective-taking spatial abilities. *Intelligence*, *32*(2), 175–191. doi: 10.1016/j.intell.2003.12.001

Hodgson, E., Bachmann, E., & Waller, D. (2011). Redirected walking to explore virtual environments: Assessing the potential for spatial interference. *ACM Transactions on Applied Perception*, *8*(4), 22:1–22:22. doi: 10.1145/2043603.2043604

Iwata, H. (1999). Walking about virtual environments on an infinite floor. *Proceedings of IEEE Virtual Reality* (pp. 286–293). Washington DC: IEEE Computer Society. doi: 101109/VR.1999.756964

Kline, P. B., & Witmer, B. G. (1996). Distance perception in virtual environments: Effects of field of view and surface texture at near distances. *Proceedings of Human Factors and Ergonomics Society 40th Annual Meeting* (Vol. 2, pp. 1112–1116). Santa Monica, CA: HFES. doi: 10.1177/154193126904002201

Knapp, J. M., & Loomis, J. M. (2004). Limited field of view of head-mounted displays is not the cause of distance underestimation in virtual environments. *Presence: Teleoperators and Virtual Environments*, *13*(5), 572–577.

Kozhevnikov, M., Motes, M. A., & Hegarty, M. (2007). Spatial visualization in physics problem solving. *Cognitive Science*, *31*(4), 549–579. doi: 10.1080/15326900701399897

Kulik, A. (2009). Building on realism and magic for designing 3D interaction techniques. *IEEE Computer Graphics and Applications*, *29*(6), 22–33. doi: 10.1109/MCG.2009.115

Loomis, J. M., Beall, A. C., Macuga, K. L., Kelly, J. W., & Smith, R. S. (2006). Visual control of action without retinal optic flow. *Psychological Science*, *17*(3), 214–221.

Loomis, J. M., & Knapp, J. M. (2003). Virtual and adaptive environments. In L. J. Hettinger & M. W. Haas (Eds.), *Virtual and adaptive environments* (pp. 21–46). Mahwah, NJ: Erlbaum.

Macuga, K. L., Loomis, J. M., Beall, A. C., & Kelly, J. W. (2006). Perception of heading without retinal optic flow. *Perception & Psychophysics*, *68*(5), 872–878.

Marsh, W. E., Kelly, J. W., Dark, V. J., & Oliver, J. H. (2012). The cognitive implications of virtual locomotion with a restricted field of view. In *Proceedings of SPIE, Volume 8289, The Engineering Reality of Virtual Reality 2012, 82890I.* doi:10.117/12.912197

Marsh, W. E., Putnam, M., Kelly, J. W., Dark, V. J., & Oliver, J. H. (2012). The cognitive implications of semi-natural virtual locomotion. *Proceedings of IEEE Virtual Reality (VR 2012: Virtual Reality Short Papers and Posters).* IEEE Computer Society.

McCreary, F. A., & Williges, R. C. (1998). Effects of age and field-of-view on spatial learning in an immersive virtual environment. *Proceedings of Human Factors and Ergonomics Society 42nd Annual Meeting* (pp. 1491–1495). Santa Monica, CA: HFES. doi: 10.1177/154193129804202106

Medina, E., Fruland, R., & Weghorst, S. (2008). Virtusphere: Walking in a human size VR "hamster ball." *Proceedings of the Human Factors and Ergonomics Society 52nd Annual Meeting* (pp. 2102–2106). Santa Monica, CA: HFES. doi: 10.1177/154193120805202704

Messing, R., & Durgin, F. H. (2005). Distance perception and the visual horizon in head-mounted displays. *ACM Transactions on Applied Perception*, *2*(3), 234–250.

Mohler, B. J., Thompson, W. B., Creem-Regehr, S. H., Pick, H., Jr., & Warren, W. H., Jr. (2007). Visual flow influences gait transition speed and preferred walking speed. *Experimental Brain Research*, *181*(2), 221–228. doi: 10.1007/S00221-007-0917-0

Montello, D. R. (1993). Scale and multiple psychologies of space. In A. Frank & I. Campari (Eds.), *Spatial information theory: A theoretical basis for GIS* (Vol. 716, pp. 312–321). Berlin: Springer-Verlag.

Péruch, P., May, M., & Wartenberg, F. (1997). Homing in virtual environments: Effects of field of view and path layout. *Perception*, *26*(3), 301–311.

Razzaque, S., Kohn, Z., & Whitton, M. C. (2001). Redirected walking. *Proceedings of Eurographics*, 289–294.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society London B*, *298*(1089), 199–209.

Still, J. D., & Dark, V. J. (2013). Cognitively describing and designing affordances. *Design Studies*, *34*(3), 285–301.

Thompson, W. B., Willemsen, P., Gooch, A. A., Creem-Regehr, S. H., Loomis, J. M., & Beall, A. C. (2004). Does the quality of the computer graphics matter when judging distances in visually immersive environments? *Presence: Teleoperators and Virtual Environments*, *13*(5), 560–571. doi: 10.1162/1054746042545292

Wandell, B. A. (1995). *Foundations of vision*. Sunderland, MA: Sinauer Associates.

Wang, Z., Bauernfeind, K., & Sugar, T. (2003). Omnidirectional treadmill system. *Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, HAPTICS 2003* (pp. 367–373). Washington DC: IEEE Computer Society. doi: 10.1109/HAPTIC.2003.1191314

Warren, W. H., Jr., Kay, B. A., Zosh, W. D., Duchon, A. P., & Sahuc, S. (2001). Optic flow is used to control human walking. *Nature Neuroscience*, *4*(2), 213–216.

Wickens, C. D., & Baker, P. (1995). Cognitive issues in virtual reality. In W. Barfield & T. A. Furness, III (Eds.), *Virtual environments and advanced interface design* (pp. 514–541). New York: Oxford University Press.

Willemsen, P., Colton, M. B., Creem-Regehr, S. H., & Thompson, W. B. (2009). The effects of head-mounted display mechanical properties and field of view on distance judgments in virtual environments. *ACM Transactions on Applied Perception*, *6*(2), 8:1–8:14. doi: 10.1145/1498700.1498702

Zanbaka, C., Lok, B., Babu, S., Ulinski, A., & Hodges, L. F. (2005). Comparison of path visualizations and cognitive measures relative to travel technique in a virtual environment. *IEEE Transactions on Visualization and Computer Graphics*, *11*(6), 694–705. doi: 10.1109/TVCG.2005.92