Bernhard E. Riecke
Max Planck Institute for Biological Cybernetics
Spemannstrasse 38
72076 Tübingen, Germany
and Vanderbilt University
Nashville, TN
bernhard.riecke@tuebingen.mpg.de

Consistent Left-Right Reversals for Visual Path Integration in Virtual Reality: More than a Failure to Update One's Heading?

Abstract

Even in state-of-the-art virtual reality (VR) setups, participants often feel lost when navigating through virtual environments. In VR applications and psychological experiments, such disorientation is often compensated for by extensive training. Here, two experimental series investigated participants' sense of direction by means of a rapid point-to-origin paradigm without any performance feedback or training. This paradigm allowed us to study participants' intuitive spatial orientation in VR while minimizing the influence of higher cognitive abilities and compensatory strategies. After visually displayed passive excursions along one- or two-segment trajectories, participants were asked to point back to the origin of locomotion "as accurately and quickly as possible." Despite using an immersive, high-quality video projection with a 84° × 63° field of view, participants' overall performance was rather poor. Moreover, about 40% of the participants exhibited striking qualitative errors, namely left-right reversals—despite not misinterpreting the visually simulated turning direction. Even when turning angles were announced in advance to obviate encoding errors due to misperceived turning angles, many participants still produced surprisingly large systematic and random errors, and perceived task difficulty and response times were unexpectedly high. Careful analysis suggests that some, but not all, of the left-right inversions can be explained by a failure to update visually displayed heading changes. Taken together, this study shows that even an immersive, high-quality video projection system is not necessarily sufficient for enabling natural and intuitive spatial orientation or automatic spatial updating in VR, even when advance information about turning angles was provided. We posit that investigating qualitative errors for basic spatial orientation tasks using, for example, rapid point-to-origin paradigms can be a powerful tool for evaluating and improving the effectiveness of VR setups in terms of enabling natural and unencumbered spatial orientation and performance. We provide some guidelines for VR system designers.

1 Introduction

Most modern virtual reality (VR) simulators suffer from a grave malady: severe disorientation (Darken & Sibert, 1996; Darken & Peterson, 2002; Grant & Magee, 1998; Lawson, Graeber, Mead, & Muth, 2002; Péruch & Gaunet, 1998; Ruddle & Jones, 2001; Ruddle, Payne, & Jones, 1998; Ruddle & Lessels, 2006). This strong tendency to easily get lost when navigating in VR can be overcome if people (a) are allowed to physically perform the simu-

Presence, *Vol. 17, No. 2, April 2008,* 143–175

lated actions (e.g., through physical walking or at least turning (Chance, Gaunet, Beall, & Loomis, 1998; Loomis, Klatzky, Golledge, & Philbeck, 1999; Klatzky, Loomis, & Golledge, 1997; Klatzky, Loomis, Beall, Chance, & Golledge, 1998; Ruddle & Lessels, 2006; Wraga, Creem-Regehr, & Proffitt, 2004), (b) are provided with useful visual landmarks or a well-known visual scene (Chance et al., 1998; Klatzky et al., 1998; Riecke, van Veen, & Bülthoff, 2002; Riecke, von der Heyde, & Bülthoff, 2005), and/or (c) are given sufficient time to employ higher cognitive processes like mental spatial reasoning and/or receive extensive feedback training on the task (Gramann, Muller, Eick, & Schonebeck, 2005; Lawton & Morrin, 1999; Riecke et al., 2002; Wiener & Mallot, 2006).

This often observed disorientation in VR stands in striking contrast to the real world, where spatial orientation and spatial updating typically operate automatically and effortlessly, requiring few if any cognitive resources (Farrell & Robertson, 1998; Loomis, Da Silva, Philbeck, & Fukusima, 1996; Presson & Montello, 1994; Rieser, 1989). Thus, most VR simulation paradigms do not empower people to use their "normal," evolutionary-developed, spatial orientation abilities. Instead, VR users often seem to resort to cognitively more demanding and computationally more expensive strategies. This might be related to the lack of robust and effortless spatial updating observed in many VR situations.

In order to determine what critical aspects of the real world are not being captured in modern VR systems, we developed an experimental paradigm that mitigates the influence of higher cognitive abilities and strategies. There are two main elements to the experimental paradigm. First, a simple and ecologically plausible task is used, rapid pointing to the origin of locomotion after visually displayed passive excursions consisting of a linear translation, a subsequent rotation, and, in some cases, a second linear translation. In a way, one could picture this task as providing the indication of a "homing vector" that points from the current position and orientation back to the starting position (Loomis et al., 1999; Klatzky et al., 1997). Rotations and translations are the basic constituents of all locomotion in the sense that even the most complex trajectories can be decomposed into a combination of elementary rotations and translations. Thus, if the most elementary combination of translations and rotations should fail, all more complex spatial orientation tasks based on path integration should also be doomed to fail. When performed in the real world using physical walking, pointing back to the origin of travel after one- or two-segment excursions is usually perceived as quite easy and not requiring much cognitive effort or computationally demanding strategies, even when performed with limited or no visual cues (Klatzky et al., 1998; Sadalla & Montello, 1989; Sholl, 1989). Using a rapid pointing paradigm has the strong advantage that it neither provides the time nor the feedback necessary to develop or use higher cognitive abilities (e.g., spatial reasoning) or strategies (Riecke, von der Heyde, & Bülthoff, 2005). It is important to note that participants in the present study never received any performance feedback. Second, by presenting only optic flow information using a uniformly textured, naturalistic ground plane, visual landmarks and other navigation aids are eliminated from the virtual environment, further restricting the possible influence of high-level strategies.

Rapid pointing after simple excursion paths is quite trivial to perform in the real world, even when all visual and auditory spatial cues and landmarks are excluded (e.g., using blindfolds and headphones displaying broadband noise). Due to an "automatic spatial updating" of our egocentric mental spatial representation of our immediate surroundings while walking, we maintain a natural and intuitive knowledge of where we are with respect to the environment during shorter periods of travel (Farrell & Robertson, 1998; Presson & Montello, 1994; Rieser, 1989). When visual and auditory cues are excluded, vestibular, proprioceptive, and kinesthetic cues are still sufficient for enabling automatic spatial updating. We may not be perfectly accurate and precise due to accumulating path integration errors during the locomotion, but the task is relatively easy to perform in the sense that it does not require noticeable cognitive effort—we just seem to automatically "know" where we are with respect to immediate objects of interest. This is typically reflected in the subjective ease of performing the task, a minimal cognitive load, a lack of qualitative errors such as left/right reversals, and rather short overall response times (typically below 2 s) with little or no dependence on the angle turned or distance traveled (Farrell & Robertson, 1998; Rieser).

When comparable tasks are performed in a virtual environment where only path-integration based visual cues (optic flow) are provided and participants are not allowed to physically move, overall response errors increase and participants typically think more before responding (Chance et al., 1998; Klatzky et al., 1998; Gramann et al., 2005; Péruch, May, & Wartenberg, 1997). For simple spatial orientation tasks like triangle completion or estimation of turning angles, both systematic and variable errors seem to depend considerably on the display device used, with head mounted displays and flat projection screens yielding the largest systematic and random errors, and large, curved projection screens yielding the lowest errors (Kearns, Warren, Duchon, & Tarr, 2002; Marlinsky, 1999; Péruch et al., 1997; Riecke, Schulte-Pelkum, & Bülthoff, 2005; Schulte-Pelkum, Riecke, von der Heyde, & Bülthoff, 2004). It seems, though, that some kind of feedback training is often critical for enabling acceptable performance in VR, even for spatial orientation tasks as simple as pointing to the origin of locomotion after short excursions. In the following, we will discuss three relevant VR-based point-to-origin studies in more detail.

1.1 Point-to-Origin Tasks in Visual VR Devoid of Any Landmarks

Lawton and Morrin (1999) displayed simple computer-simulated rectangular mazes on a desktop monitor and asked participants to point back to the origin of travel after excursions of 3, 5, or 7 segments using a compass-like pointer. Despite the simple geometry of the maze and path layout (constant segment lengths with 90° in between turns), pointing performance showed considerable errors even for the simplest condition: Mean absolute pointing errors averaged around 40° for men and 60° for women and increased for increasing number of path segments. Participants who maintained some kind of "feeling" of the relative direction of the origin (similar to a homing vector) performed significantly better than those who did not. Conversely, remembering the sequence of left and right turns had detrimental effects on pointing accuracy, suggesting that more cognitive strategies based on route knowledge cannot necessarily compensate for the apparent lack of natural, intuitive spatial orientation in VR. When participants were repeatedly asked to indicate the homing vector during the excursion, however, pointing errors decreased by about 10° for both men and women. Providing pointing feedback only at the end of the excursion did not improve pointing performance significantly, though. The data suggests that continuously maintaining a representation of the direction toward the origin of travel (similar to a homing vector) was critical for good pointing performance at the end of the trajectory.

In a recent point-to-origin task performed in desktop VR (Gramann et al., 2005), participants followed a visually displayed uniformly textured tunnel consisting of straight and curved segments, and were asked at the end of the excursion to indicate the direction to the origin of travel (homing vector) by adjusting a simulated 3D arrow using mouse buttons. Participants were given repeated feedback about the correct pointing direction, which might have contributed to the relatively low absolute pointing errors (10–25°). Differences between initial and final heading never exceeded 60°, which largely reduces the range of sensible pointing directions and might also have contributed to the good overall performance. To obviate this limitation, the current experiment was designed to maximize the range of correct pointing directions to span the whole range from small angles (as low as ±10°) to large angles (±180°). Furthermore, the experiments described in this paper also recorded response times and used a more immersive, projection-based VR system with a larger FOV.

In order to investigate the influence of path complexity on visual path integration performance, Wiener and Mallot (2006) used a joystick-based point-to-origin paradigm in a simple virtual environment consisting of a uniformly textured ground plane presented on a flat back-projection screen (90° × 60° FOV). Given sufficient feedback during an initial training phase, participants were able to perform the purely visual point-to-origin tasks with reasonable accuracy (20–35° absolute pointing error), even when the excursion path included up to 4 turns (albeit always in the same direction). Performance was moreover independent of the number of turns. Response times were, however, always above two seconds, suggesting that the task was not perceived as simple. This was corroborated by subjective reports of participants and the amount of errors during the training phase. That is, instead of using quick and robust automatic spatial updating as in the real world, participants apparently had to resort to different, computationally more demanding strategies.

The three above-mentioned VR-based studies all used extensive feedback training and unlimited response times. This allowed for fairly accurate pointing performance. In the present study, however, we aimed at investigating how well participants perform when they are *never* provided with any performance feedback and are asked to respond as "accurately and quickly as possible"—factors that are critical for the overall acceptance and usability of VR.

1.2 Left-Right Errors and the Apparent Failure to Update Rotations That Are Not Physically Performed

There is an increasing amount of research demonstrating that spatial perception in VR is prone to systematic errors such as misestimation of distances or turning angles (see, e.g., Riecke, Schulte-Pelkum, & Bülthoff, 2005; Thompson et al., 2004, and references therein). Apart from those *quantitative* errors, there are also a few accounts of *qualitative* errors that cannot be simply explained by a systematic misperception of distances traveled or angles turned.

Klatzky et al. (1998) were the first to report an apparent failure to update heading changes that were not physically performed. When participants were asked to imagine walking along a verbally described two-segment excursion and respond by turning to face the origin as if they had actually walked the trajectory, participants responded as if standing at the to-be-imagined location, but still facing the initial orientation. This resulted in qualitative errors, namely left-right errors. When the excursion path contained, for example, a leftward turn, the proper turn-to-face-origin response would have been to also turn leftward by less than 180°. Instead, though, participants turned rightward, thus producing a left-right error.[1] Presenting only optic flow information via a head mounted display in a control condition resulted in similar qualitative errors and apparent failures to update the visually presented turn. Only when participants actually walked the path or at least physically executed the turn between the two segments did they properly incorporate the rotation, which corroborates the often posited importance of physical motion cues for automatic spatial updating (e.g., Farrell & Robertson, 1998; May & Klatzky, 2000; Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004).

More recently, systematic left-right inversions have also been observed in desktop VR experiments by Gramann et al. (2005). In a 30-trial categorization pre-test, participants saw visually simulated passive excursions along simple curved tunnels. After the excursion, two arrows (one facing left, the other one facing right) appeared on the computer monitor, and participants were asked to select the one that pointed to the origin of locomotion. In the main test, participants used mouse buttons to rotate a visually displayed arrow such that it pointed to the origin of locomotion. In total, 23 of 43 participants (the so-called non-turners) responded as if they had not updated their heading and were still facing the original orientation, thus producing left-right inversions. Gramann et al. argued that the non-turners in their study used an allocentric strategy, whereas the turners—who did incorporate the heading changes—used an egocentric strategy.

Personal communications with the authors of Wiener and Mallot (2006) revealed that some participants in their point-to-origin study initially produced left-right errors as well. Over the course of the feedback training phase, however, those left-right errors quickly disap-

1. Note that we use the term "left-right error" or "left-right inversion" in a purely descriptive sense, without any implication about the underlying processes that might have caused the left-right errors. Those underlying processes might, for example, include failures to update rotations, actual left-right confusion, or consistently choosing an ineffective strategy.

peared. Apart from methodological differences, one noteworthy difference between the turn-to-face-origin study by Klatzky et al. (1998) and the VR studies by Gramann et al. (2005) and Wiener and Mallot (2006) is that all participants in the former study showed left-right errors, whereas only a subset of the participants in the latter studies showed such qualitative errors. Note that participants in the Klatzky et al. study performed only five trials each, and it is conceivable that extended exposure might have led them to realize their mistake and adjust their behavior.

The current study was designed to investigate the striking phenomenon of left-right inversions in more detail by using a wide range of turning angles (30–170°) and a large number of trials per person. Most importantly, we never provided any performance feedback or training that participants could have used to correct for potential errors. In particular, and as an important difference to the procedures of Gramann et al. (2005), we did not provide turners and non-turners with different, strategy-specific performance feedback to reinforce their strategy. Furthermore, we did not exclude any participants because they switched between turner and non-turner strategies as was done in the Gramann et al. study. Instead, we performed a real-world practice phase where blindfolded participants walked along several two-segment paths and used the same point-to-origin paradigm as in the main experiment to ensure that they clearly understood the task requirements and procedure and knew how to use the pointing device. In addition to the two-segment experiments, we performed a subsequent one-segment experiment where participants simply had to point back after a visually depicted translation, followed by a rotation about varying angles. Using a turn-to-face-origin procedure and a verbal description of the outbound trajectory similar to Klatzky et al. (1998), Avraamides, Klatzky, Loomis, and Golledge (2004) showed that left-right errors vanished completely for two subsequent one-segment trials. Here, we wanted to test whether potential left-right errors in VR would also disappear for one-segment trials. Avraamides et al. argued that the left-right errors observed for the two-segment task were caused by a failure to update the cognitive heading. Furthermore, they argued that participants noticed their mistake in the subsequent one-segment task and corrected for them, as not updating one's heading would invariably have led to the same response—namely, turning 180°, irrespective of the turning angle. Thus, if potential left-right errors in our study were caused by a failure to update rotations similar to Avraamides et al., those left-right errors should be expected to disappear for our one-segment experiment.

1.3 Reference Frame Conflict and Presence in VR

Whenever a stimulus other than the real world is used in experiments, participants are confronted by two, possibly interfering, representations of the environment (May, 2004; Riecke & von der Heyde, 2002; Riecke & McNamara, 2007; Wang, 2005). On the one hand, the representation of the physical surround (e.g., the physical VR setup and the surrounding lab space); and on the other hand, the representation of the simulated or intended scene (presented typically on a visual display). According to a theoretical spatial orientation framework by Riecke and von der Heyde (2002), the ability of VR users to quickly and intuitively orient themselves while navigating should be dependent on the degree to which they feel spatially present[2] in the simulated scene. Presence, in turn, should be impaired by the interference between the two representations or reference frames. Thus, the model predicts that both presence and quick and intuitive spatial orientation should be impaired if the participant experiences a conflict between the simulated motion through the virtual environment and the stationary physical surround of the actual room. Hence, at least parts of the observed difficulty in spatially updating visually simulated ego-motions might be due to the conflict between the (intended) representation of the simulated space and the (ideally to-be-ignored) representation of the physical surround (Riecke, Cunningham, & Bülthoff, 2007). As a step toward testing these

2. Presence is here conceptualized as the "subjective experience of being in one place or environment, even when one is physically situated in another" (Witmer & Singer, 1998). See also IJsselsteijn (2004) and Sadowski and Stanney (2002) for recent reviews on presence and related issues.

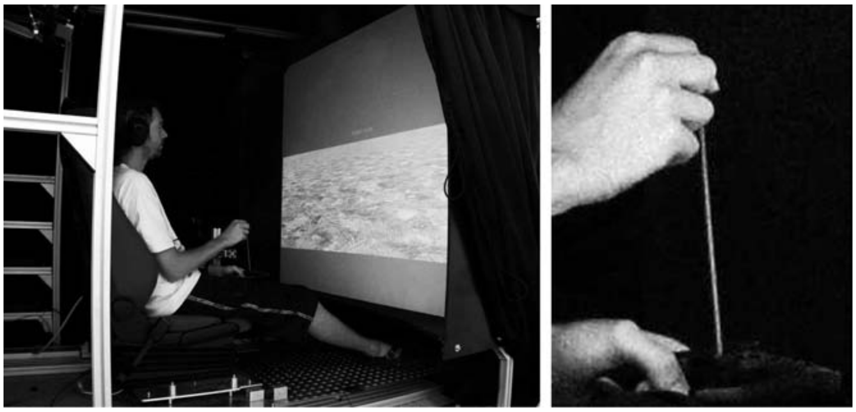

Figure 1. *VR system showing a participant with the pointing device (modified gamepad) seated in front of the projection screen displaying the textured ground plane devoid of any landmarks.*

predictions, we compared two conditions where participants could either see and hear the physical surrounding lab (low immersion condition) or not (high immersion condition) in a first experimental series. A second experimental series investigated whether spatial orientation performance in optic flow-based VR can be improved by providing explicit advance knowledge of turning angles.

2 Experimental Series 1

2.1 Methods

Sixteen naive participants (gender-balanced, aged 13–39 years, with a mean age of 23.75 years) completed the first experimental series.[3] Participation was voluntary and paid at standard rates. All participants had normal or corrected-to-normal vision.

2.1.1 Stimuli and Apparatus. Participants were seated at a distance of 89 cm from a flat projection screen (1.68 m width × 1.26 m height, corresponding to a field of view of about 84° × 63°), as illustrated in Figure 1. The virtual environment was quite simple and consisted of a textured flat ground plane that did not contain any absolute orientation or distance cues. We chose a texture that mimics a grass-like surface to reduce the artificial and unnatural appearance often associated with optic flow displays. The ground plane texture was designed to contain both a broad range of spatial frequencies and a high contrast in order to provide strong optic flow cues about the distance traveled and angles turned. Note, however, that the virtual scene did not contain any useful landmark information that participants could have used for determining their position or orientation relative to the origin of locomotion. Visual stimuli were projected non-stereoscopically using a JVC D-ILA DLA-SX21S video projector with a resolution of 1400 × 1050 pixels.

In the high immersion condition, participants wore active noise canceling headphones (Sennheiser HMEC 300) playing several mixed layers of flowing water to exclude all external noise. In addition, the curtains on both sides of the projection screen were closed, such that participants could neither see nor hear the surrounding laboratory (see Figure 1). In a second, low immersion condition, participants wore no headphones, and the curtains on both sides of the projection screen were opened, such that the surrounding lab was visible.

3. The first experimental series is in part based on a conference paper by Riecke and Wiener (2007).

Care was taken to adjust the light level in the lab to be similar to the curtains in the high immersion condition. We hypothesized that the high-immersion condition might help to reduce the conflict between the simulated virtual scene (depicting a simulated self-motion) and the real world (i.e., the VR setup and surrounding lab, which was stationary; Riecke et al., 2007; Riecke & McNamara, 2007), and thus indirectly facilitate spatial orientation relative to the virtual scene rather than the real world (see also Prothero, 1998).

2.1.2 Procedure. Each trial consisted of a passive motion phase, a pointing phase, and a fixed inter-trial interval. The motion phase consisted of a translation along a first segment s_1 (8 m/s maximum translational velocity, with a brief acceleration and deceleration phase to avoid motion sickness), followed by a rotation (30 deg/s) on the spot, and a subsequent translation along a second segment s_2 (same velocity as for s_1). For the one-segment experiments, the second translation was omitted. Upon arriving at the end of the trajectory, participants were asked to point "as accurately and quickly as possible" to the origin of locomotion as if they had physically moved (pointing phase). The inter-trial interval consisted of a 3 s period where the screen was blanked, and a 2 s interval where participants were instructed to prepare for the next trial. The turning direction was alternated between trials to reduce the occurrence of potential motion aftereffects and motion sickness, but was not analyzed separately.

Pointing was performed using a modified game pad where the knob was replaced by an 18 cm long thin (2 mm) plastic rod to allow for more precise pointings (see Figure 1). Participants were instructed to hold the top end of the rod with the index finger and thumb of their preferred hand. The direction of rod deflection indicated the pointing direction, and a pointing was recorded once the joystick was deflected by more than 95%. Pre-tests had shown that this allows for more accurate pointing than simply using a joystick (which is often used in pointing studies), most likely because one uses a precision grip on a long, straight rod that is rotationally symmetric.

Compared to (real or simulated) compass-like pointers that are sometimes used for point-to-origin experiments (Gramann et al., 2005; Lawton & Morrin, 1999; Muehl & Sholl, 2004; Sadalla & Montello, 1989; Sholl, 1989), using a rapid pointing paradigm with an upright default position of the pointer (e.g., Riecke, von der Heyde, & Bülthoff, 2005) has the advantage of allowing for equally short response times for all pointing directions. Furthermore, the upright default position ensures that there is no directional bias and participants have (from a bio-mechanical perspective) similar pointing motions and response times for all directions, a problem that is often not accounted for in studies using compass-like pointers.

Participants indicated that the pointing device was easy and intuitive to use. Note that participants were never provided with any performance feedback throughout the experiment to mitigate the usage of cognitive strategies or recalibration. Furthermore, participants were asked to point "as accurately and quickly as possible" to reduce the likelihood of their building up any abstract geometric representations—for example, a top-down view of the path geometry, as was observed in experiments where participants were given unlimited response time (Riecke et al., 2002). This was important for the purpose of the experiment, as we were interested in testing whether participants were able to orient themselves naturally, that is, quickly and intuitively (most likely through automatic spatial updating), without any need for feedback training and/or computationally expensive processing. Previous studies had shown that participants can indeed perform triangle completion and point-to-origin tasks in VR relatively well if given unlimited response time and sufficient feedback training (Gramann et al., 2005; Lawton & Morrin, 1999; Wiener & Mallot, 2006).

2.2 Experimental Design

The experimental series consisted of the following parts: a demonstration phase, followed by a real-world practice phase, a two-segment familiarization experiment, a two-segment main experiment, a one-segment experiment, and a post-experimental debriefing.

2.2.0.1 Demonstration Phase. To become familiar with the experimental task and procedures, participants were given a few demonstration trials by the experimenter and received written and oral instructions.

2.2.0.2 Real-World Practice Phase. Participants were asked to walk physically with eyes closed along five or more two-segment paths in the actual lab and use the pointing device (which was for that purpose detached from the computer) to point back to the origin of locomotion. Pointing back to the origin of locomotion after a two-segment real-world excursion was perceived as rather trivial, but served well to familiarize participants with the experimental task and pointing device without providing them with any specific feedback that could be used in the actual VR experiment. In fact, none of the participants showed any problems or qualitative errors (like left/right confusion) in the practice phase from the very beginning, and quantitative errors were minimal, suggesting that the pointing device and procedure introduced little, if any, systematic or random errors. Once participants indicated that they did not need any more practice trials and clearly understood the instructions, experimental procedures, and task requirements for the VR test, they proceeded with the familiarization experiment. For all VR conditions, participants were instructed to treat the visual motion simulation as if it originated from an actual self-motion, and to respond as if they had actually moved (just like in the real-world practice phase).

2.2.0.3 Two-Segment Familiarization Experiment. In order to reduce the impact of learning effects on the main experiment, all participants first performed a familiarization experiment. The familiarization experiment consisted of 22 trials, each consisting of a factorial combination of two lengths of s_1 (16 m, 24 m) × five turning angles γ (45°, 75°, 105°, 135°, 165°) × two turning directions (left, right; alternating), plus four additional baseline trials without any rotation ($\gamma = 0°$, four lengths of s_1 (16 m, 24 m) × two repetitions). The familiarization experiment lasted about 10 minutes on average. The turning angles were selected to be different from those used in the main experiment in order to avoid potential direct learning transfer or memorization of turning angles. For each participant, the immersion condition for the familiarization experiment matched that of the first session of the main experiment.

2.2.0.4 Two-Segment Main Experiment. After completing the practice phase and familiarization experiment, participants performed a two-segment main experiment which was split into two sessions (high immersion condition and low immersion condition) in balanced order. Each of the two sessions of the main experiment was composed of 52 trials, consisting of a factorial combination of two lengths of s_1 (16 m, 24 m; randomized) × six turning angles γ (30°, 60°, 90°, 120°, 150°, 170°; randomized) × two turning directions (left, right; alternating) × two repetitions per condition, plus four baseline trials (randomly interspersed) without any turns between the two segments (two lengths of s_1 (16 m, 24 m; randomized) × two repetitions for $\gamma = 0°$).

2.2.0.5 One-Segment Experiment. A large amount of variation across different studies seems to be caused by problems in perceiving and encoding visually simulated turns (Klatzky et al., 1998; Chance et al., 1998; Riecke et al., 2002; Riecke, Schulte-Pelkum, & Bülthoff, 2005). This naturally raises the question as to whether some of the errors observed in the main experiment above are caused by problems in veridically perceiving and encoding the visually presented turning angles. To control for this possibility, and to test whether potential left-right errors would disappear for the one-segment task as predicted by Avraamides et al. (2004), all participants performed a subsequent one-segment experiment. The task was simply to point back to the origin of locomotion after being presented with a visually simulated passive forward translation ($s_1 = 16$ m) followed by a passive rotation with angle γ, but no additional second translation. As in the main experiment, each participant performed two sessions (high immersion and low immersion) in balanced order (same order as before). The one-segment experiment consisted of 28 trials per session: a factorial combination of six turning angles γ (30°, 60°, 90°, 120°, 150°, 170°; randomized) × two turning directions (left, right; alternating) × two repetitions per condition, plus four baseline trials (randomly interspersed) without any turn after the translation (four repetitions for $\gamma = 0°$).

2.2.0.6 One-Segment Encoding Control Experiment. In order to rule out the possibility that the new pointing device induced a systematic measurement error and to further investigate the potential influence of misperceiving the visually displayed rotations, we ran a new set of seven naive but psychophysically experienced observers (lab members, all male) in a modified version of the one-segment experiment (low immersion condition). Unlike in the previous experiment, participants were now given explicit advance information about the upcoming turn. That is, participants were told verbally about the exact turning angle and turning direction (e.g., "120° left") prior to the onset of each trial (and thus, in principle, had all the information they needed to determine the location of the origin). This procedure should essentially eliminate all errors from the encoding phase (building up an internal representation of the angle turned and trajectory traveled), such that all remaining errors should stem from problems with determining the proper response (mental spatial reasoning phase) and/or problems in actually performing the intended pointing response (execution phase). See Riecke et al. (2002) for a discussion of these three different phases in the context of a triangle completion task in VR.

2.2.1 Dependent Variables. Pointing performance was analyzed in terms of three dependent variables. The response time was calculated as the time until the pointer was deflected by 95%. The absolute pointing error per participant was computed as the mean absolute value of the difference between the correct homing direction and pointing direction indicated by the participant per trial. Instead of analyzing the signed pointing error or bias—which is problematic with pointing data due to their 360° periodicity—we used circular statistics to compute mean pointing directions (Batschelet, 1981) and performed a more graphical data analysis (see Figure 2, 3, 4, 5, and 7, which will be discussed later), which we hope will improve the understandability and interpretability of the data. In addition, we computed the circular standard deviation, which can be conceived as the circular statistics counterpart of the standard deviation and is a measure of the variability or consistency of the pointing data per participant and condition (Batschelet, 1981).

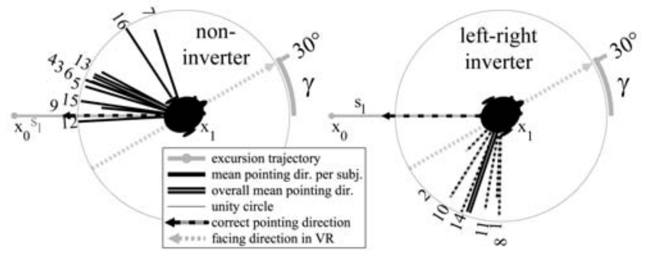

Figure 2. *Sample data for the 30° condition of the one-segment experiment, split up into the six participants who showed systematic left-right errors (right subplot) and the remaining ten participants who did not show such systematic left-right errors (left subplot). Plotted is a top-down schematic view of the excursion path (in solid gray) from the start point x_0 to the endpoint x_1 and the subsequent turn by 30°. The mean pointing direction of each participant is indicated by the different bars and subject IDs. The length of the mean pointing vector indicates the consistency of the individual pointing directions: Shorter mean pointing vectors indicate higher circular standard deviations of the individual pointing (e.g., participant 2), whereas mean pointing vectors close to the surrounding black unity circle indicate high consistency and thus low circular standard deviations of the individual pointings (e.g., participant 14 and 16; Batschelet, 1981).*

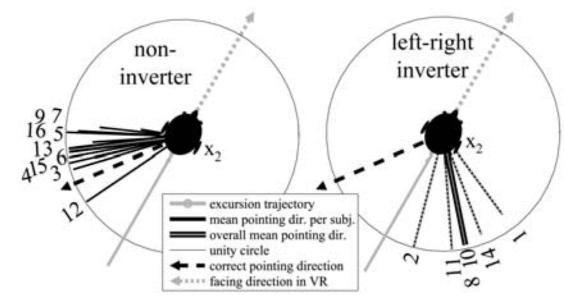

Figure 3. *Sample data for the 60° condition of the two-segment main experiment illustrating the systematic left-right errors for 6 of the 16 participants (right subplot). The data are plotted as in Figure 2 and represents a close-up of the endpoint of the 60° trajectories (cf. Figure 4).*

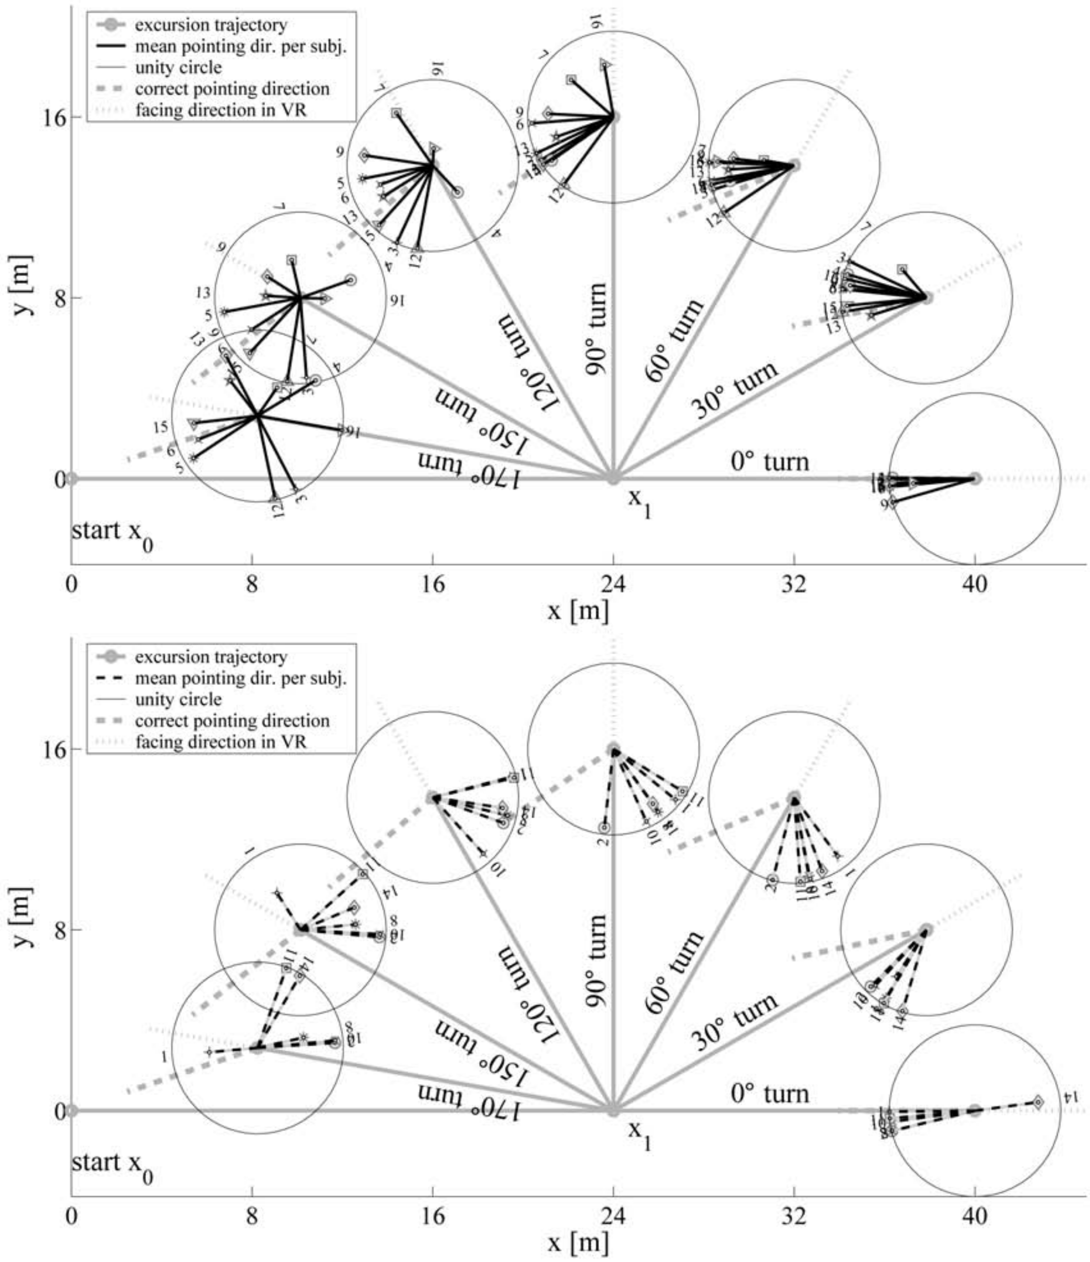

Figure 4. *Mean pointing directions for s_1 = 24 m and the different turning angles of the two-segment main experiment, plotted as in Figure 3. The bottom subplots represent data from the six left-right inverters (depicted with dashed lines), the top subplots show data from the ten non-inverters (solid lines). Corresponding data plots for s_1 = 16 m are available at http://www.kyb.mpg.de/publication.html?publ=4781.*

2.2.1.1 Mental Spatial Abilities Test. To investigate potential relations between spatial orientation performance as assessed by the rapid pointing paradigm and more general spatial abilities, participants were asked to perform a standard paper-and-pencil mental spatial abilities test (Stumpf & Fay, 1983) after the VR

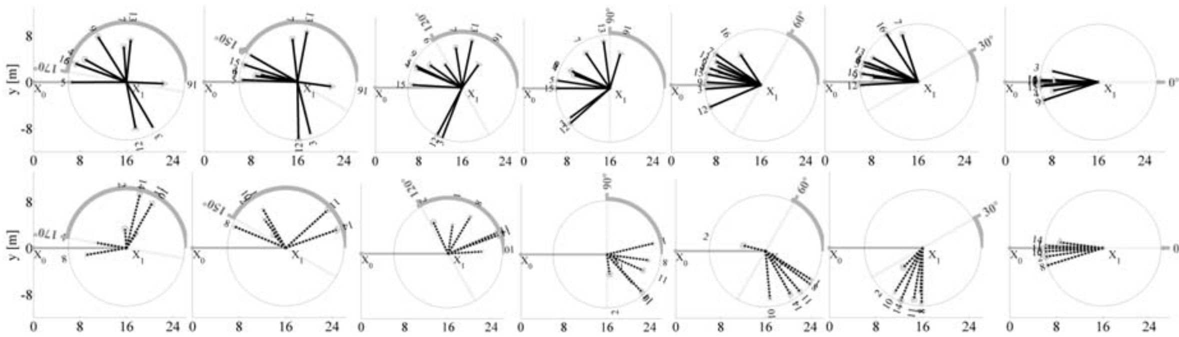

Figure 5. *Mean pointing directions for the different turning angles γ of the one-segment experiment, plotted as in Figure 2 and separated into left-right inverters (bottom) and non-inverters (top). Note the increasing absolute pointing errors and within- and between-subject pointing variability for increasing turning angles.*

experiment. In the spatial abilities tests, participants saw for each of the 21 trials a picture of a curved tube located within a transparent quadratic box, and had to judge in a multiple-choice manner from which viewpoint a second picture of the same cube was taken.

2.3 Results and Discussion

The pointing data were quantified using repeated measures ANOVAs for all three dependent measures and the factors immersion condition, length of the first segment s_1, and turning angle γ. The no-turn condition ($\gamma = 0°$) was excluded from the ANOVA as it was intended as a baseline condition. Surprisingly, the immersion condition did not show any significant main effects at all for any of the dependent measures. Thus, it seems as if presence and immersion did not play an important role for the point-to-origin task used, and/or the manipulation was too subtle to be effective. For the further analysis, the data were pooled over the two immersion conditions and the turning directions (which were not the focus of the current study). The pooled data are summarized in Figures 2, 3, 4, 5, and 6.[4]

2.3.1 Pointing Errors. As can be seen in Figures 2, 3, 4, and 5, the pointing data were rather noisy and showed considerable variability both within-subjects and between-subjects.

2.3.1.1 Consistent Left-Right Inversions. The pointing data showed a bimodal distribution of the participant population with respect to their pointing behavior. This is nicely illustrated for the pointing responses for 30° rotations in Figure 2. Ten of the 16 participants pointed leftward (left subplot), which is at least roughly the direction toward the origin, whereas the pointing directions for the other six participants (right subplot) seem to be mirrored with respect to the current observer orientation in the virtual scene. Careful analysis of all the experimental conditions in Figure 7 revealed that the participant population clustered indeed into two distinct groups that exhibited qualitatively different overall pointing behavior. For turns to the left, the proper pointing direction is always to the left, and vice versa. Ten of the 16 participants indeed pointed consistently in the correct overall direction, that is, leftward for left turns and rightward for right turns (at least for turning angles $\gamma \leq 90°$). The other six participants pointed, however, consistently in the wrong direction (see Figures 2–4). That is, when the excursion path contained a counterclockwise (left) turn, they pointed consistently to the right instead of to the left, and vice versa, even though left turns should always result in leftward pointings for turning angles $<180°$. This group of participants will in the following be termed "left-right

4. Note that high-resolution color versions of all figures of this paper are available at http://www.kyb.mpg.de/publication.html?publ=4781.

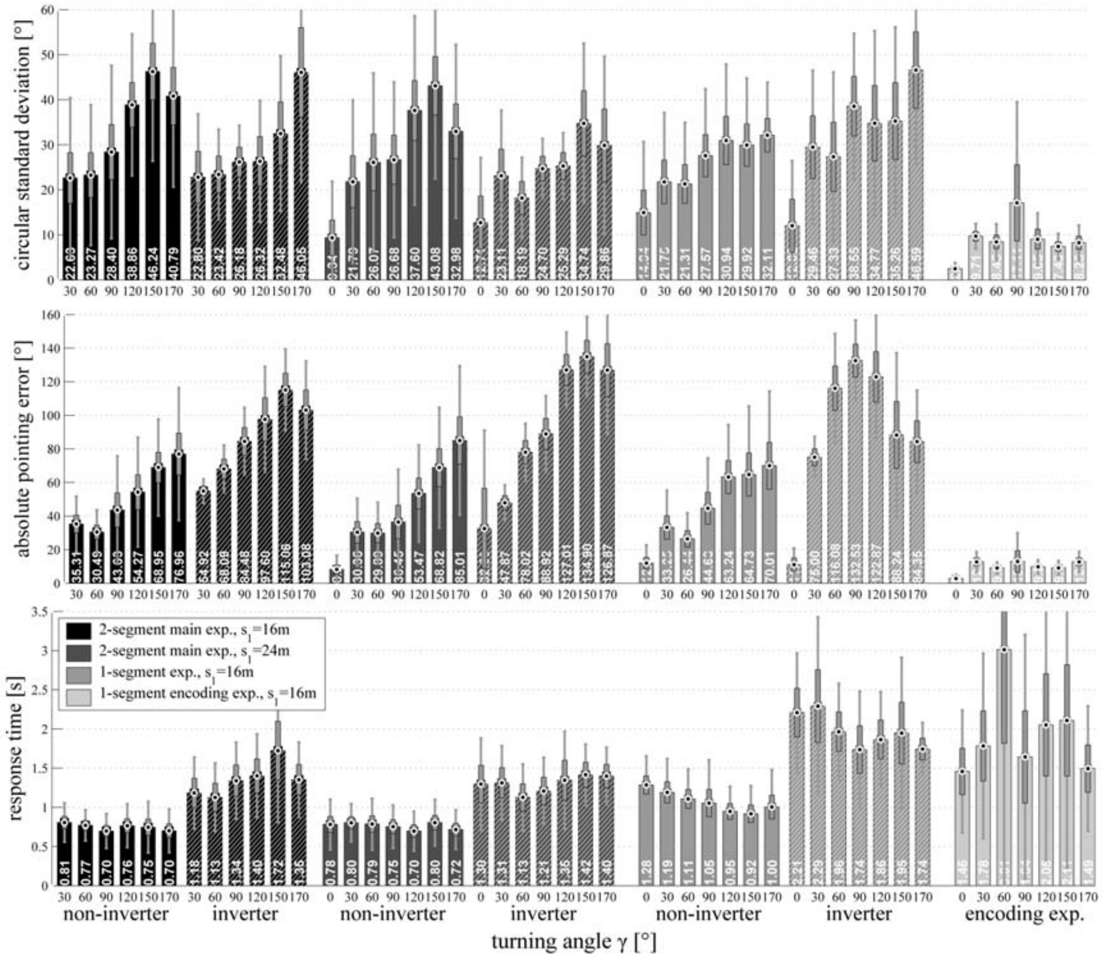

Figure 6. *Summary of the arithmetic means of the circular standard deviation (top), absolute pointing error (middle), and response time (bottom). Solid and hatched bars represent data from the non-inverters and inverters, respectively. Boxes and whiskers indicate one standard error of the mean and one standard deviation, respectively.*

inverters." Note that such left-right errors are to the best of our knowledge not known from blindfolded walking studies (see introduction and general discussion). These left-right errors are most clearly visible for smaller turning angles. For larger turns, pointing directions are more noisy and left-right side errors might be confounded with the large misestimations of the actual turning angle[5] indicated in Figure 5. If the presented turning angles are overestimated, a 150° left turn might, for example, be perceived as a 200° left turn, and the

5. As an attempt to resolve the 360° ambiguity of the pointing data, the following algorithm was used for the data plotting and gain factor analysis of Figures 7 and 12: For the one-segment data, the mean pointing directions per participant were plotted step-by-step for increasing turning angles (starting with 0°). When the difference between the mean pointing direction for the current turning angle (e.g., 90°) and the next larger one (here: 120°) differed by more than 180°, the mean pointing direction of the latter was remapped to an interval of ±180° surrounding the former (using a modulo 360° operation). The overall good linear fits in Figure 7 suggest that this procedure was successful, as good linear fits indicate a rather constant overestimation or underestimation of the turning angles for each of the participants—which is in agreement with previous results on turn estimation using a similar VR setup (Riecke, Schulte-Pelkum, & Bülthoff, 2005). A similar algorithm was used for the two-segment data in Figures 7 and 12 for decreasing correct egocentric homing directions.

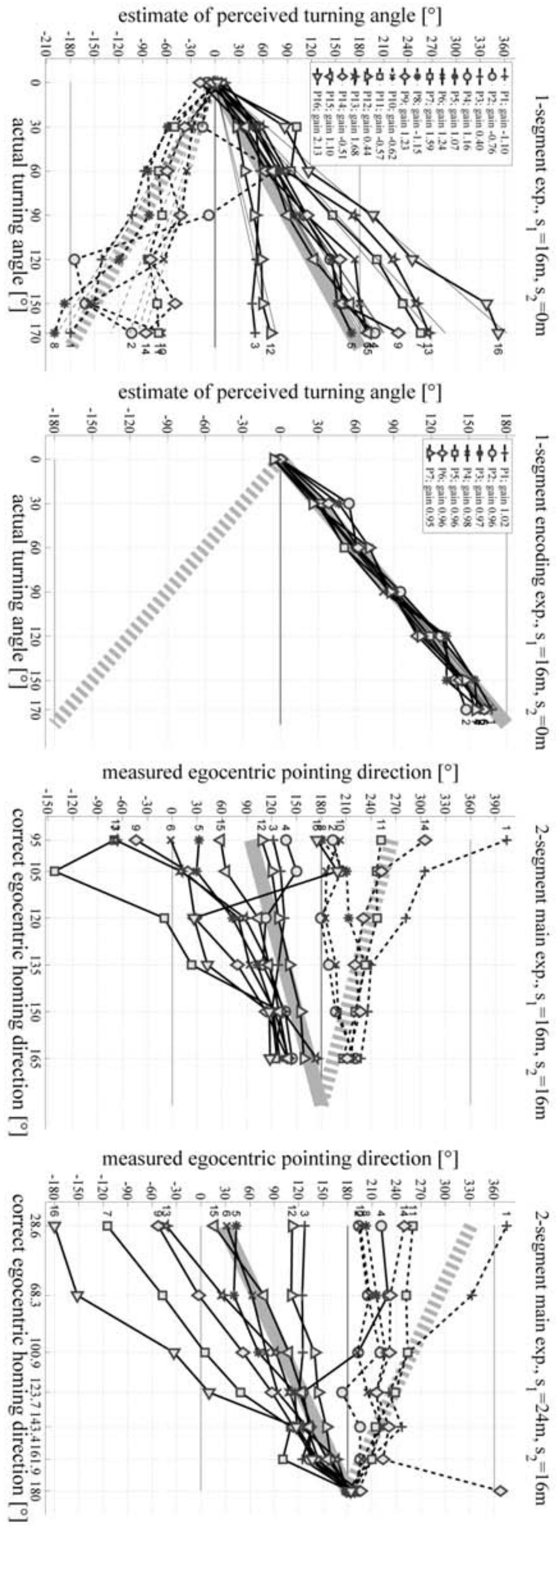

Figure 7. *Left plots: For the one-segment experiments, an estimate of the perceived turning angle was computed by taking 180° minus the measured egocentric pointing direction. That is, the estimate of the perceived turning angle was defined as the turning angle that would correspond to a given egocentric pointing direction if the encoding of the traveled trajectory and the mental computation and execution of the pointing response were all free of errors. Thus, negative values of the estimated perceived turning angle indicate only that participants responded as if they mistook left turns for right turns and vice versa, even though there are, of course, other underlying processes that might produce similar data. The thick gray diagonal line indicates the expected response for perfect performance; the thick dashed gray line denotes the expected response for consistent left-right swap errors. Linear least squares fits were used to compute the slope or gain factor between the estimated perceived turning angle and the actual turning angle for the different participants, and are indicated in the top inset of the top figures. This plotting method shows a bimodal distribution of gain factors, and was in fact used to categorize participants as non-inverters or inverters: Participants who showed a positive slope and values predominately above 0 were categorized as non-inverters (plotted with solid lines), whereas participants with negative slopes and values predominately below 0 were categorized as left-right inverters (plotted with dashed lines). The second subplot from the left shows data from the encoding control experiment, where participants were explicitly told the turning angle before each trial. Right plots: Participants' mean egocentric pointing direction, plotted over the correct egocentric pointing direction (i.e., the homing direction). Note that the above procedure of estimating perceived turning angles was only applicable for the one-segment data and not the two-segment data, as the two-segment task required considerably more complex mental spatial reasoning, and the estimate of the perceived turning angle for the two-segment data would have been confounded with the perception of the traveled distance—and the data suggests that some participants might not have been able to clearly distinguish between the two values of s_1. Note the overall large errors for all but a few participants. As for the one-segment graphs, this plotting method illustrates nicely the difference between the non-inverters (solid lines), which tend to have an overall positive slope in their response and values below 180°, and the left-right inverters (dashed lines), who showed an overall flat or negative slope and values predominately above 180°, indicating that they pointed into the overall wrong hemisphere (i.e., rightward for left turns and vice versa).*

Table 1. *Analysis of Variance Results for the Circular* SD *(Top), Absolute Pointing Error (Middle), and Response Time (Bottom)†*

	Two-segment experiment						One-segment experiment					
	Left-right inversion		Turning angle γ		Interaction LR inversion γ		Left-right inversion		Turning angle γ		Interaction LR inversion γ	
	$F(1,14)$	p	$F(5,70)$	p	$F(5,70)$	p	$F(1,14)$	p	$F(5,70)$	p	$F(5,70)$	p
Circular *SD*	0.451	.513	**10.0**	**<.005*****	1.40	.23	1.42	.25	**3.83**	**.004****	0.51	0.77
Absolute pointing error	**18.6**	**.001*****	**36.2**	**<.005*****	**3.63**	**.006***	**17.7**	**.001****	**5.14**	**<.0005*****	**6.92**	**<.0005*****
Response time	**12.6**	**.003****	1.68	.15	**3.20**	**.012***	**12.5**	**.003****	*2.23*	*0.61 m*	0.86	.51

†For the two-segment experiment, the length of the first segment s_1 did not show any significant effects or interactions and was thus excluded from the table. Significant and marginally significant effects are typeset in bold and italics, respectively.
*$\alpha = 5\%$.
**$\alpha = 1\%$.
***$\alpha = 0.1\%$.

resulting pointing direction would then be rightward (consistent with a 200° turn, as for participant ID 13) and not leftward (consistent with a 150° turn).

To account for the bimodal distribution of the participant population, a new set of repeated measures ANOVAs were performed with the within-subject factors length of the first segment s_1 and turning angle γ and the additional between-subject factor of left-right inversion. As before, the no-turn condition ($\gamma = 0°$) was excluded from the ANOVA as it was intended as a baseline condition. The results of the statistical analysis are summarized in Table 1 and discussed below.

2.3.1.2 Absolute Pointing Error. Figure 6 shows a strong increase in absolute pointing errors for increasing turning angles γ, which is corroborated by the ANOVA results in Table 1. For the two-segment main experiment, the mean absolute pointing error increased from 30.4° for the 30° turns to 83.0° for the 170° turns. The one-segment experiment showed a similar increase from 32.3° to 72.4° for increasing turning angles. The factor s_1 did not show any significant main effects or interactions.

As expected, overall absolute pointing errors were significantly higher for the left-right inverters as compared to the non-inverters (93.9° vs. 51.1°, respectively, for the two-segment experiment and 103.2° vs. 50.4°, respectively, for the one-segment experiment, see also Figure 6 and Table 1). The significant interactions between left-right inversion and turning angle indicates a differential influence of the turning angle for the inverters versus non-inverters.

2.3.1.3 Circular Standard Deviation. The circular standard deviation (which is also referred to as the mean angular deviation) is a circular statistics measure of the within-subject variability or inconsistency in the pointing directions for the different conditions (similar to the standard deviation in linear statistics) and showed a clear influence of the turning angle γ for both the two- and one-segment experiment (see Table 1). As can be seen in Figure 6, the circular standard deviation was lowest in the no-turn condition and gradually increased with increasing turning angles. This implies that participants were less consistent (more variable) in their pointing responses for increasing turning angles, which might be related to the increased task difficulty for larger turning angles and/or increasing uncertainties in the estimation of the turning angles.

Inverters and non-inverters did not differ significantly in their pointing consistency and were equally affected by the turning angle manipulation, indicated by the lack of any significant effects or interactions for the factor left-right inversion (see Table 1).

2.3.2 Response Times. Mean response times in the familiarization experiment were already quite short (1.31 s, with 1.46 s *SD*) and decreased further in the main experiment (0.97 s on average, 0.50 s *SD;* see Figure 6). Response times for the two-segment experiment did not show any significant relation to the turning angle (see Table 1). Unexpectedly, response times in the one-segment experiment were 1.37 s on average (*SD:* 0.80 s) and thus noticeably higher than in the main experiment. As more processing time was apparently needed directly after a rotation, one might argue that participants might have perceived the rotations as more difficult to update and needed more processing time for rotations than translations. This bears some resemblance to spatial updating studies where rotations are typically found to be harder to imagine than translations (Rieser, 1989; May, 1996; Presson & Montello, 1994). Interestingly, the one-segment experiment revealed a marginally significant tendency toward decreasing response times for larger turning angles ($F(5, 70) = 2.23$, $p = .061$; cf. Table 1), as if turns became easier to update with increasing turning angles. In particular, the smallest turning angle (30°) resulted in the highest mean response time (1.60 s), compared to the largest turning angle of 170° (1.28 s response time). This is the opposite of what one might have expected from studies on mental rotations or imagined self-rotations, where response times typically increase with increasing rotation angle (Rieser; Shepard & Metzler, 1971). This puzzling result cannot be convincingly explained by the current data and awaits further investigation. One might, for example, speculate that the processing of the turning direction (left or right) might have increased participants' processing time specifically for the smaller turning angles, where they had to respond shortly after perceiving the turning direction.

Response times were considerably longer for the inverters as compared to the non-inverters (1.33 s vs. 0.75 s, respectively, for the two-segment experiment and 1.92 s vs. 1.04 s, respectively, for the one-segment experiment), suggesting that the inverters might have needed more cognitive resources to perform the task.

2.3.3 Correlation Between Left-Right Inversion and Post-Experimental Data. Even though the reasons underlying the observed left-right inversions are not fully understood yet (see general discussion), it is interesting to note that left-right inversion was associated with lower spatial abilities as measured using a standard paper-and-pencil test (Stumpf & Fay, 1983; 8.7 points for the inverters vs. 15.2 for the non-inverters; $t(14) = 4.54$, $p < .0005$***; see Figure 8). Furthermore, it turned out that five of the six left-right inverters were female, whereas the non-inverters were predominately male (7/10, $\chi^2(1, N = 16) = 4.27$, $p = .039$*). Despite gender being known to correlate with many spatial abilities (see Coluccia & Louse, 2004, and Lawton & Morrin, 1999, for comprehensive reviews), the observed gender bias should be interpreted with care, as many other factors apart from gender might contribute to spatial abilities. For example, none of the inverters had any 3D computer game experience, whereas 6 of the 10 non-inverters did ($\chi^2(1, N = 16) = 5.76$, $p = .016$*). Furthermore, the number of participants in the current study was simply too small to allow for comprehensive analysis of gender effects. Note that general computer usage did not show a significant difference between the two groups (1.52 vs. 2.36 h/day on average for the inverters and non-inverters, respectively; $t(14) = 1.14$, $p = .27$; see Figure 8).

2.3.3.1 Behavioral Versus Verbal Ratings of Turning Angles. Figures 4, 5, and 7 reveal a general tendency to overestimate turning angles for the non-inverters and underestimate turning angles for the inverters. This can be quantified by the mean gain factors (defined as the slopes in Figure 7, left subplot), which were 1.20 (*SD* = 0.53) for the non-inverters and −0.79 (*SD* = 0.28) for the inverters. When participants were asked to verbally rate the maximum turning angle in a post-experimental interview, however, a different data pattern was observed. Estimates for the maximum

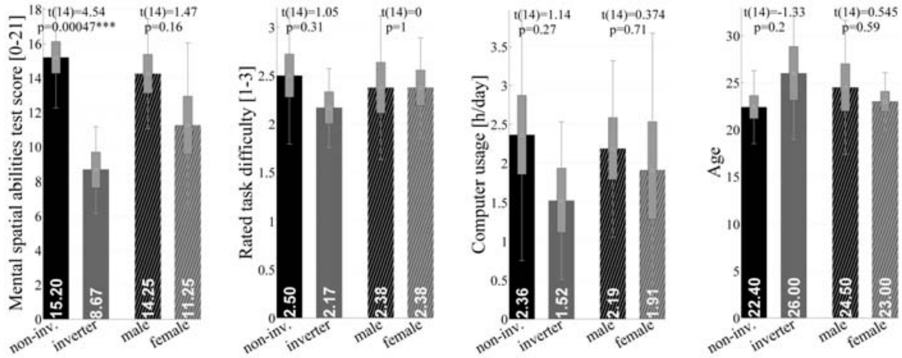

Figure 8. *Data from the post-experimental questionnaire for the main experiment. Bars, boxes, and whiskers indicate the arithmetic mean, one standard error of the mean, and one standard deviation, respectively. At the top of each plot, between-subject (unpaired) t-tests indicate pairwise comparison for left-right inversion (left) and gender (right).*

turning angle ranged from 120–450° (M = 283°). That is, participants overestimated the maximum turning angles by 66% on average and 165% maximally, corresponding to gain factors of 1.66 and 2.65, respectively. That is, the overestimation in the verbal ratings were much higher than the overestimation computed from the behavioral data. Furthermore, the verbal data did not show any difference between the inverters and non-inverters. Ratings for the maximum turning angle were on average 300° for the inverters and 273° for the non-inverters. This suggests different processes underlying the verbal and behavioral responses.

2.3.3.2 Rated FOV. The horizontal FOV subtended by the projection screen was estimated as 100° on average, with values ranging from 60–150°. This corresponds to a slight overestimation by 19% on average, and is surprisingly similar to the behavioral measure of the overestimation of the turning angles for the non-inverters (1.20, see previous paragraph), suggesting that the overestimation of the FOV might at least in part be related and contributing to the apparent overestimation of turning angles. The general overestimation of visually displayed rotations was rather unexpected, as an earlier study using the same setup showed instead a small underestimation of visual turns that were actively controlled (Schulte-Pelkum et al., 2004). The underlying reasons are not understood yet. Hence, a given VR system should be carefully evaluated for a given stimulus and task if veridical perception of simulated self-rotations is intended (see also Riecke, Schulte-Pelkum, & Bülthoff, 2005).

2.3.3.3 Vection. It is interesting to note that the duration of the visual motions was much shorter than the average onset latencies of vection (i.e., the time it takes until observers experience a visually-induced self-motion illusion). We had earlier measured average vection onset latencies of about 15 s for linear forward motions and about 12 s for rotational motions for the same setup (Riecke, Schulte-Pelkum, Caniard & Bülthoff, 2005; Riecke, Schulte-Pelkum, & Caniard, 2006). Informal reports of participants confirm this: None of the observers had any convincing sense of self-motion in any of the trials during the experiment. This leaves open the question whether a compelling feeling of self-motion might have enhanced spatial updating performance in the current study and might thus help to improve spatial orientation in VR in general. We are currently running an experiment to test this hypothesis.

2.3.4 One-Segment Encoding Control Experiment. Data from the one-segment encoding control experiment is summarized in Figures 9, 6, and 7. Com-

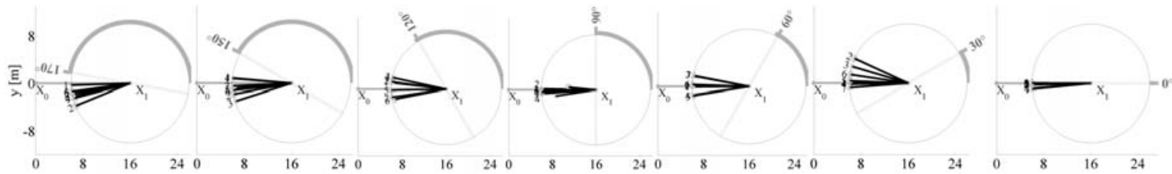

Figure 9. *Mean pointing directions for the one-segment encoding control experiment, plotted as in Figure 5. Note the high pointing accuracy (low systematic pointing errors) and precision (low within- and between-subject variability, indicated by the mean pointing vector length close to 1).*

pared to the previous one-segment experiment, the encoding experiment where participants had advance knowledge about the upcoming turning angle and direction showed a considerable decrease in pointing errors and pointing variability (both within- and between-subjects, see Figure 7). In fact, systematic pointing errors were minimal. This suggests that the large pointing errors found in the previous experiments might be largely caused by an encoding error or a mental computation error. That is, participants might have misperceived the turning angle, and/or might have been unable to infer or mentally compute the proper response from the given information. Note, however, that differences in the participant population might also have contributed. While participants for the main experiment were psychophysically inexperienced observers, participants for the encoding control experiment were lab members and thus psychophysically experienced observers. The second experimental series was designed to address this issue explicitly.

Despite using experienced observers, though, overall response times for the encoding control experiment were somewhat higher than in the previous experiments, which suggests that participants needed more cognitive resources to perform the task. This is consistent with participants stating that they perceived the task as "unexpectedly extremely difficult"—despite having extensive experience with VR simulations.

The high overall pointing accuracy and low circular standard deviation in the encoding control experiment suggests that the pointing device and procedure induced little systematic execution errors. This is consistent with observations from the real-world practice phase, where participants showed little systematic or random errors. In fact, despite using a rather simple pointing device, response times, pointing accuracy, and circular standard deviations in the encoding control experiment were roughly comparable to previous spatial updating experiments that used a similar rapid pointing paradigm but a technically more advanced, two-handed, position-tracked pointing wand (Riecke, von der Heyde, & Bülthoff, 2005; Riecke et al., 2007). Furthermore, all participants of the current study reported that the pointing device was easy and intuitive to use. Taken together, this suggests that the pointing methodology was appropriate and did not contribute to the qualitative errors and the high perceived difficulty of the point-to-origin task.

In summary, it seems as if providing explicit advance information about the exact turning angle might be a way to obviate encoding errors and systematic misperceptions of turning angles and in turn enable excellent performance for simple optic flow-based spatial orientation tasks in VR. If this were indeed the case, it would not only be interesting from the pure science perspective, but also from an applied perspective of effective-yet-affordable self-motion simulation.

The second experimental series was designed to test the potential benefit of advance information about turning angles more thoroughly. Note that only experienced psychophysical observers (seven lab members who were naive to the specific task but had extensive experience with VR experiments) participated in the one-segment encoding control experiment, which might have contributed to the low systematic and random errors. To test whether these results generalize, the second experimental series used a new set of 24 naive participants who had no extensive experience with VR experiments and were not lab members. Furthermore, we aimed at replicating and extending the previous findings about the striking phenomenon of left-right inversion and the observed correlations between left-right inversion and subject-specific factors like general spatial abilities, 3D computer game experience, and gender. To maximize comparability, the second experimental series closely replicated the procedures of the first experimental series, apart from providing additional advance information about the upcoming turning angle and using a more thorough post-experimental debriefing.

3 Experimental Series 2

3.1 Methods

A new set of 24 naive participants (gender-balanced) completed the second experimental series.[6] All participants had normal or corrected-to-normal vision. Participants' age ranged from 21 to 41 years (M = 26.4, *SD:* 5.4). Participation was voluntary and paid at standard rates. Apart from the differences described below, stimuli, apparatus, experimental procedures, data analysis, and plotting closely matched those of the first experimental series in order to allow for direct comparisons.

While the announcement of the upcoming turning angle had been given verbally for the previous one-segment control experiment, turning angles were announced visually for the second experimental series. That is, during the 4 s instruction phase, the turning angle of the upcoming trajectory was displayed at the center of the screen (e.g., "120°"). Note that the turning direction was not indicated, as participants did not seem to have any problems determining the proper turning direction from the presented visual motion, despite the occurrence of left-right inversions.

6. Two additional participants had to be excluded as they had clearly misunderstood the experimental instructions, as became obvious in the post-experimental debriefing.

As in the encoding control experiment of the previous experimental series, the goal of providing advance information about the upcoming turning angle was to reduce the influence of systematic or random errors due to a potential misperception in the visually presented turning angles. This, in turn, was expected to mitigate errors during the encoding phase (building up an internal representation of the angle turned and trajectory traveled), such that all remaining errors should stem from problems with determining the proper response (mental spatial reasoning phase) and/or problems in actually performing the intended pointing response (execution phase). See Riecke et al. (2002) for a discussion of these three different phases in the context of a triangle completion task in VR. As the high- and low-immersion condition in the previous experiments did not show any significant difference, the second experimental series was only run in the high-immersion condition. To allow for a more comprehensive understanding of the data, the post-experimental debriefing was extended. Apart from these differences, the second experimental series was identical to the first experimental series and consisted of the same parts: a demonstration phase, followed by a real-world practice phase, a two-segment familiarization experiment, a two-segment main experiment, a one-segment experiment, and a post-experimental debriefing.

3.2 Results and Discussion

The pointing data were pooled over left and right turns and are graphically represented in Figures 10, 11, 12, and 13.

3.2.0.1 Consistent Left-Right Inversions. As can be seen in Figures 10, 11, and 12, the participant population showed as before a bimodal distribution and clustered into 11 left-right inverters and 13 non-inverters: When presented with, for example, trajectories containing left turns, the left-right inverters pointed predominately to the right instead of to the left. This is nicely illustrated in Figure 12, where the inverters' pointing responses are overall closer to the left-right mirrored response (gray dashed line) than the correct response

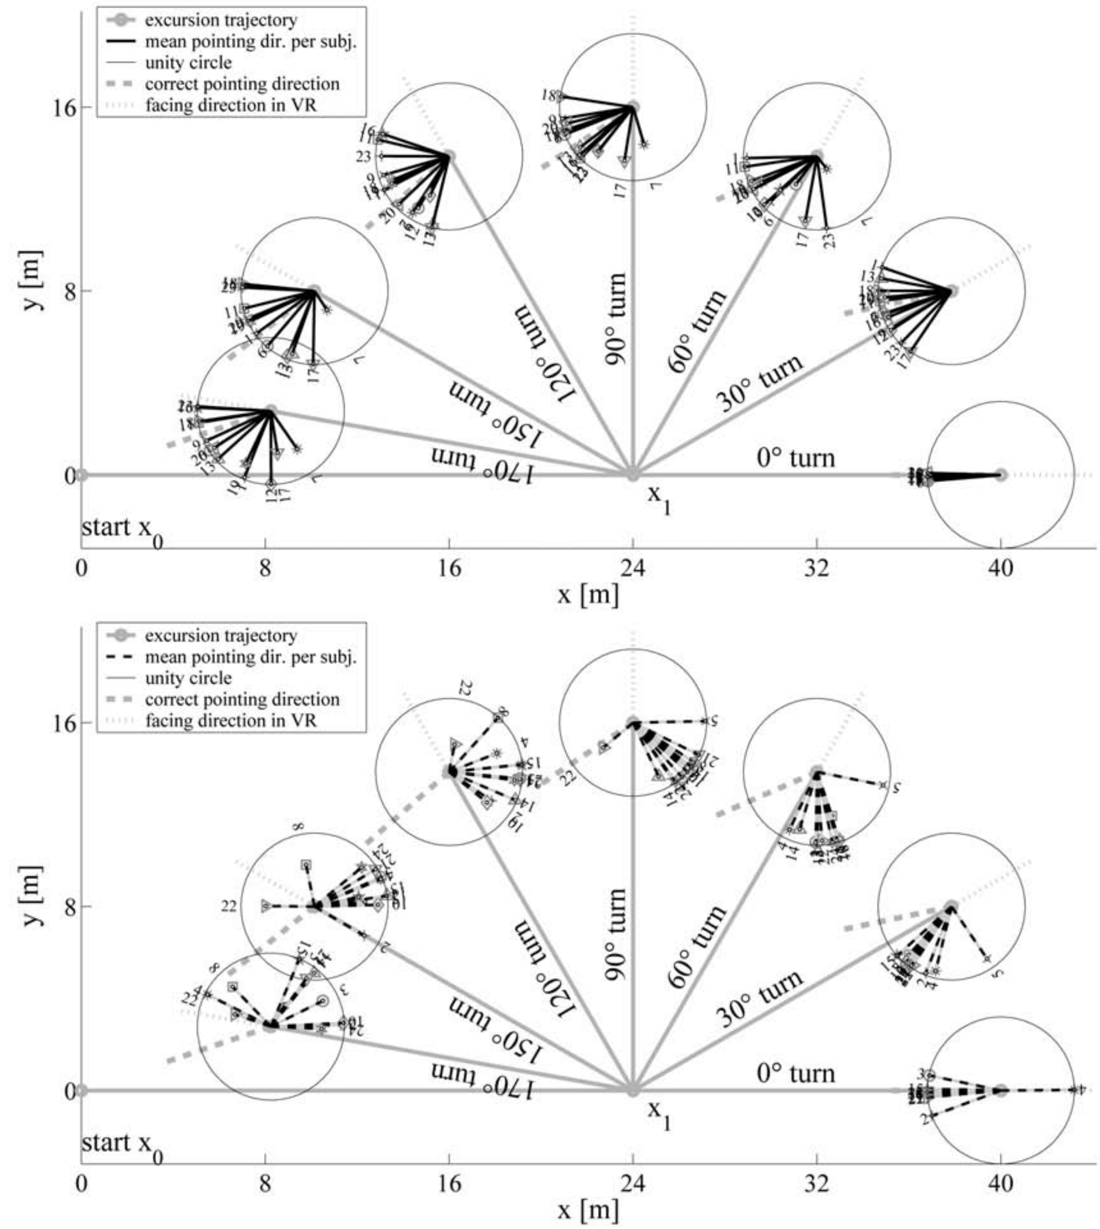

Figure 10. *Top-down schematic view of the two-segment data for s_1 = 24 m from the second experimental series, plotted as in Figure 4. The bottom subplot represents data from the eleven left-right inverters (depicted with dashed lines); the top subplot shows data from the 13 non-inverters (solid lines). Note the surprisingly high between-subject variability despite the advance information about the turning angle. Corresponding data plots for s_1 = 16 m are available at http://www.kyb.mpg.de/publication.html?publ=4781.*

(solid gray line). Together with the bimodal distribution, this was taken as the criterion for categorizing participants as left-right inverters. This occurrence of consistent left-right inversions corroborates the findings of the first experimental series, and the ratio of left-right inverters in the current study was even slightly higher

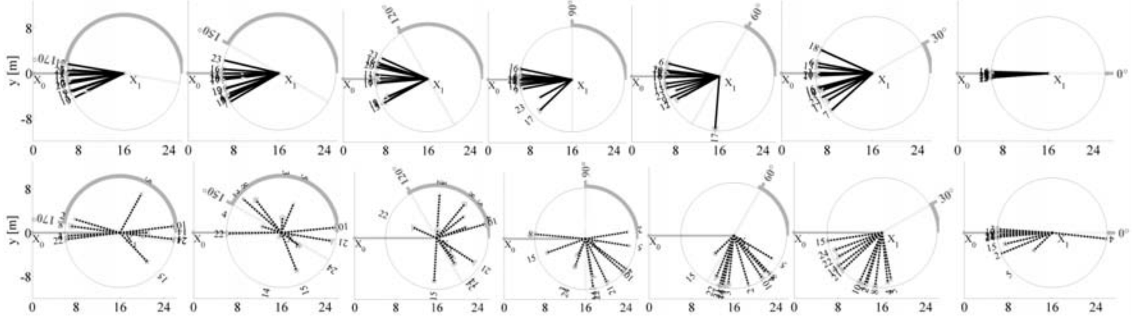

Figure 11. *Top-down schematic view of the one-segment data from the second experimental series, separated into inverters (bottom) and non-inverters (top) and plotted as in Figure 5. Note that the within- and between-subject pointing variability increased for increasing turning angles for the left-right inverters, but not the non-inverters.*

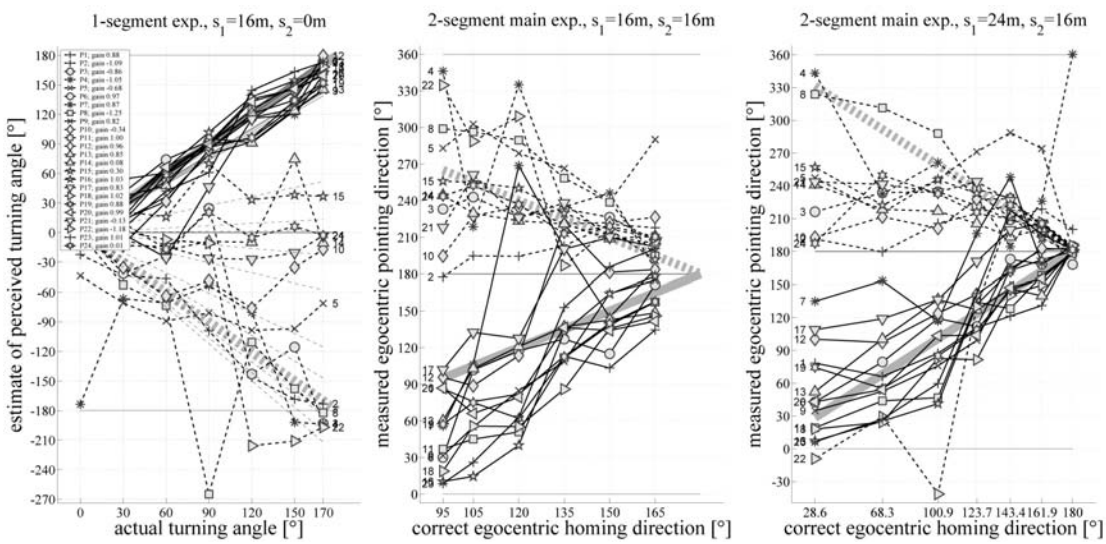

Figure 12. *Graphical analysis of the pointing data for the one-segment experiment (left) and two-segment experiment (middle and right) of the second experimental series, plotted as in Figure 7. The bimodal distribution of the participants' pointing behavior into 11 left-right inverters (dashed lines) and 13 non-inverters (solid lines) is clearly visible, although the group of left-right inverters unexpectedly seem to have a larger between-subject variability than the inverters of the first experimental series. Note the high overall errors for the two-segment task, even for the non-inverters.*

(11 of 24 participants or 46%) than in the previous study (6 of 16 or 38%). This suggests that left-right inversion in path integration-based VR is indeed a reliable and reproducible phenomenon and not an artifact that might have been caused by random peculiarities in the participant population used or by using desktop VR as in Gramann et al. (2005).

3.2.0.2 Large Pointing Variability and Errors Despite Advance Information about Turning Angles. Despite providing advance information about the turning angles, the pointing data were still surprisingly noisy and showed a considerable within- and between-subject variability, even in the seemingly simple one-segment experiment (see Figures 12 and 13). This is in striking

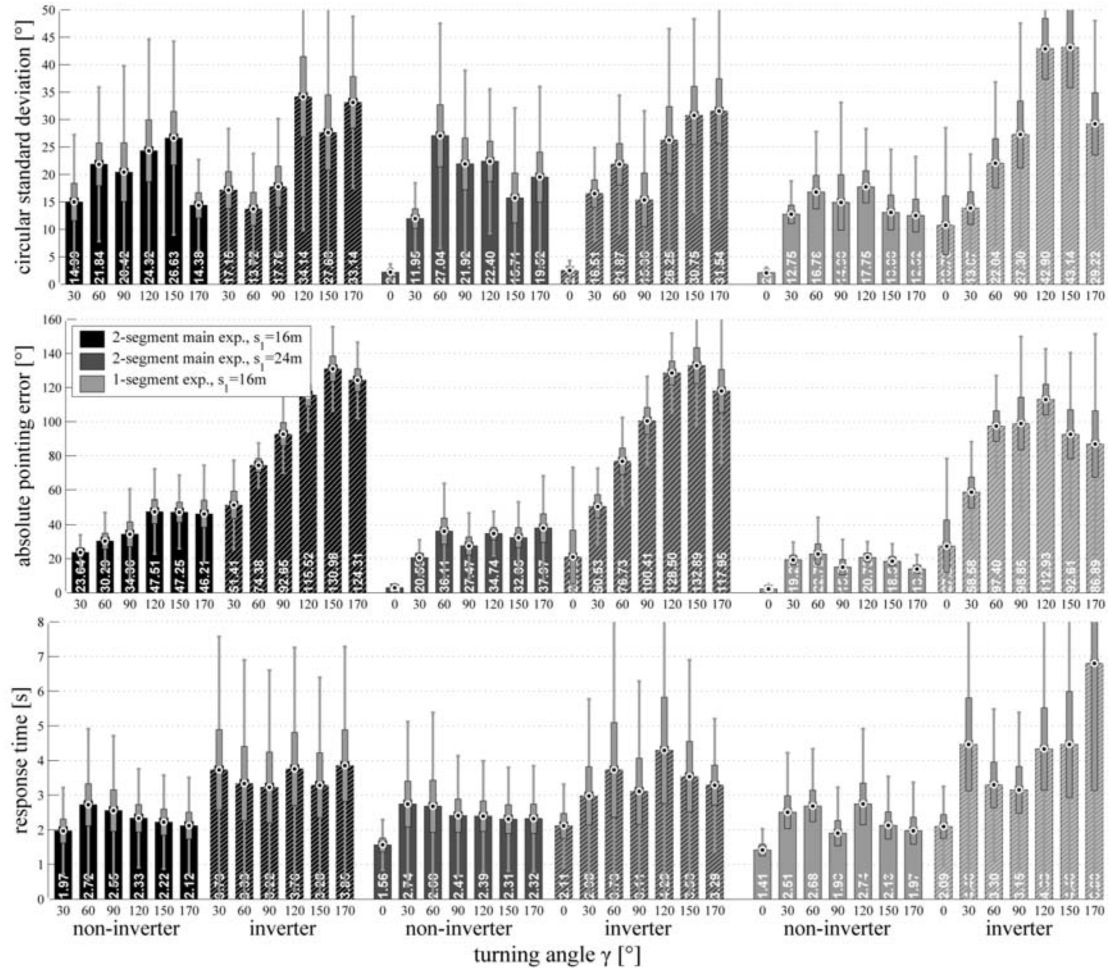

Figure 13. *Summary of the arithmetic means of the circular standard deviation (top), absolute pointing error (middle), and response time (bottom). Solid and hatched bars represent data from the non-inverters and inverters, respectively. Boxes and whiskers indicate one standard error of the mean and one standard deviation, respectively. Note that for the one-segment data, both the circular standard deviation and the absolute pointing error of the left-right inverters showed a clear increase for increasing turning angles, whereas the non-inverters showed no such effect and overall lower values.*

contrast to the corresponding encoding control experiment in the first experimental series, where participants performed with almost negligible systematic and random errors (see Figures 6 and 7).

3.2.1 Quantitative Data Analysis (ANOVAs). To provide a more quantitative analysis of the data, separate repeated measures ANOVAs were conducted for the three dependent measures response time, circular standard deviation, and the absolute pointing error with the factors length of the first segment s_1, turning angle γ, and left-right inversion. The no-turn condition ($\gamma = 0°$) was excluded from the ANOVA and the further discussion of the data as it was intended as a baseline condition. ANOVA results are summarized in Table 2 and will be discussed in the following.

3.2.1.1 Response Times. Mean response times were 2.90 s for the two-segment experiment and differed considerably between participants, indicated by the large

Table 2. *Analysis of Variance Results for the Circular Standard Deviation (Top) and Absolute Pointing Error (Bottom) of the Second Experimental Series†*

	Two-segment experiment						One-segment experiment					
	Left-right inversion		Turning angle γ		Interaction LR inversion γ		Left-right inversion		Turning angle γ		Interaction LR inversion γ	
	$F(1,22)$	p	$F(5,110)$	p	$F(5,110)$	p	$F(1,22)$	p	$F(5,110)$	p	$F(5,110)$	p
Circular *SD*	0.995	0.33	3.65	.004**	3.44	**.006***	17.46	<.0005***	4.38	**.001***	4.13	**.002***
Absolute pointing error	166.3	<.0005***	33.71	<.0005***	14.89	<.0005***	85.89	<.0005***	2.50	**.035***	2.62	**.028***

†The response time data is not listed as it did not show any significant effects or interactions. Similarly, for the two-segment experiment, the length of the first segment s_1 did not show any significant effects or interactions and was thus excluded from the ANOVA table.
*α = 5%.
**α = 1%.
***α = 0.1%.

standard deviation of 2.71 s (see also Figure 13). The one-segment experiment showed a similarly long response time (3.28 s) and an even more pronounced between-subject variability (*SD:* 4.30 s). Response times showed a trend toward higher values for the left-right inverters, both for the two-segment experiment (3.51 s vs. 2.40 s for the non-inverters) and the one-segment experiment (4.42 s vs. 2.32 s). This trend did not reach significance, though ($p > .1$). Similarly, neither turning angle nor length of s_1 or any of the interactions show any significant effects on response times in the ANOVA (see also Figure 13).

The current response times were on average more than 1 s longer than in the one-segment encoding control experiment of the first experimental series (mean: 2.02 s, *SD:* 1.82 s), which used psychophysically experienced observers who had extended experience with VR experiments. Furthermore, the current response times were much higher and more variable than those in the main experiments of the first experimental series that used a similar, naive participant population as in the current study, but provided no advance feedback about turning angles (mean response time of 0.97 s, *SD:* 0.50 s for the two-segment experiment and 1.37 s, *SD:* 0.80 s for the one-segment experiment; see Figure 6). This seems to indicate that the point-to-origin task does not become easier or require less cognitive resources when advance information about the upcoming turning angle is provided, which is in agreement with participants' post-experimental reports (see Figure 14).

3.2.1.2 Absolute Pointing Error and Circular Standard Deviation. Both the absolute pointing error and the circular standard deviation showed a clear increase with increasing turning angles for both the two- and one-segment experiments (see Figure 13 and Table 2). This was somewhat unexpected, given that participants had advance information about the exact turning angle per trial. Hence, it seems unlikely that the observed increase in pointing errors and inconsistency was caused by an increase in the uncertainty about the angle turned (which was announced beforehand). Instead, it might have been caused by an increase in task difficulty and cognitive requirements for the larger rotations. That is, one might argue that the larger rotations were harder to update than the smaller ones, similar to findings from mental rotation tasks (e.g., Shepard & Metzler, 1971). Note, however, that larger rotations did

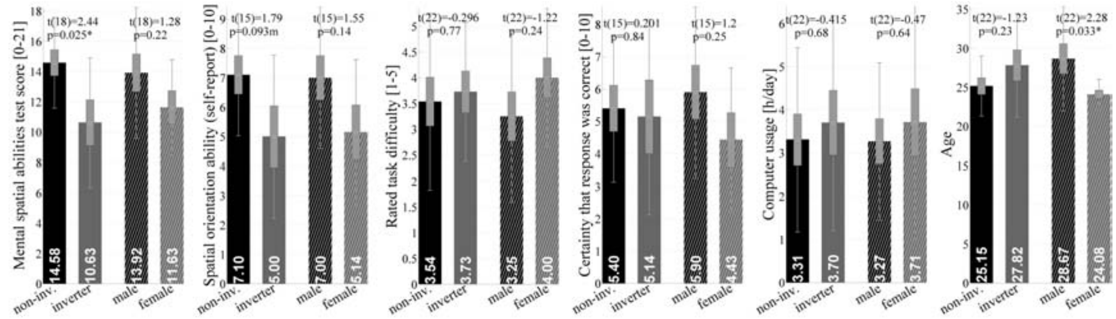

Figure 14. *Data from the post-experimental questionnaire of the experiment series 2. Bars, boxes, and whiskers indicate the arithmetic mean, one standard error of the mean, and one standard deviation, respectively. At the top of each plot, between-subject (unpaired) t-tests indicate pairwise comparison for left-right inversion (left, solid bars) and gender (right, hatched bars).*

not yield increased response times, which is somewhat puzzling.

3.2.1.3 Correlation between Left-Right Inversion and Pointing Data. Left-right inverters produced overall larger absolute pointing errors, both for the two-segment experiment (99.7° vs. 34.9° for the non-inverters) and the one-segment experiment (91.2° vs. 18.4°), as supported by the ANOVA results in Table 2. While this increase in absolute pointing error for the two-segment experiment might be simply caused by the left-right inversion, this cannot explain the observed increase for the one-segment experiment, where mere left-right inversion would yield identical absolute errors. Hence, it seems as if the inverters performed more poorly *in addition* to the left-right inversion, potentially because they did not just point into the left-right mirrored direction but also produced some other kind of systematic or random error. This is consistent with the observed significant interaction between left-right inversion and turning angle (cf. Table 2). Furthermore, left-right inverters produced significantly larger circular standard deviations in the one-segment experiment, as compared to the non-inverters (29.7° vs. 14.6°, see Table 2). This further corroborates the above conjecture. Note, though, that even the non-inverters produced slightly higher absolute pointing errors (18.4°) and circular standard deviations (14.6°) than the psychophysically experienced participants in the one-segment encoding control experiment of the first experimental series (11.1° and 7.4°, respectively). This performance advantage of the experienced observers is also well visible in Figure 12.

The one-segment experiment showed an interesting interaction between left-right inversion and the turning angle (see Table 2 and Figure 13). While the inverters showed increasing circular standard deviations and absolute pointing errors for increasing turning angles, the non-inverters showed no such influence and overall lower values. This suggests an increasing task difficulty and uncertainty for the inverters, even though one might argue that the one-segment task should be rather trivial, given the advance information about the turning angle and the relatively simple required mental computation (e.g., egocentric pointing direction = 180° – instructed turning angle, in the same direction as the simulated self-motion).

3.2.2 Correlation between Left-Right Inversion and Post-Experimental Data. *3.2.2.1 Spatial Abilities.* Left-right inverters scored significantly lower on a standard mental spatial abilities questionnaire[7] by Stumpf and Fay (1983) than the non-inverters (cf. Fig-

7. Due to technical problems, the data from four participants was missing for the mental spatial abilities test.

ure 14), thus corroborating the findings from the first experimental series. In addition to the spatial abilities test, participants were asked to rate their general spatial orientation ability in the real world on a scale from 0 (poor) to 10 (excellent).[8] Non-inverters reported slightly higher spatial orientation ability in the self-reports (see Figure 14), but this trend reached only marginal significance.

3.2.2.2. Rated Task Difficulty and Certainty that One's Response Was Correct. On average, the pointing task was perceived as relatively difficult (3.63 on a scale from 1 (very easy) to 5 (very difficult), *SD:* 1.5, see Figure 14). In fact, 11 of the 24 participants rated the task as very difficult. When asked to rate how certain they were that their pointing response was correct, participants averaged 5.3 points (*SD:* 2.5) on a scale from 0 (very uncertain) to 10 (very certain). Furthermore, only one participant reported being quite certain that his response was correct (score ≥ 9), whereas three participants reported being quite uncertain (score ≤ 2). This corroborates the surprisingly high difficulty of the point-to-origin task and suggests that participants did, in fact, not have any intuitive feeling of where they were with respect to the origin, even though the outbound paths consisted only of one- or two-segment trajectories. In fact, when explicitly asked, not a single participant reported having a real-world-like, intuitive spatial orientation. Unexpectedly, neither left-right inversion nor gender showed any systematic differences in terms of perceived task difficulty or rated certainty that the given response was correct (see Figure 14).

3.2.2.3 Perceived Path Layout. In order to assess the perceived layout of the path, all participants performed two more two-segment trials after the experiment and were debriefed directly after each trial. The first trial was a 60° left turn with $s_1 = s_2$; the second trial was a 120° right turn with $s_1 = 1.5s_2$. When asked to report the turning angle (in degrees), 16 of the 24 participants (66.7%) correctly responded for the 60° rotation. Responses ranged from 30° to 90°. The 120° turn was correctly identified by only 9 of the 24 participants (37.5%), with responses ranging between 60° and 170°. The frequent misestimation of turning angles was somewhat unexpected, as participants had extensive experience (93 trials) from the previous two- and one-segment experiments where they had explicit knowledge about the upcoming turning angle. This corroborates the general difficulty in assessing visually displayed turning angles when using a flat projection screen (Riecke, Schulte-Pelkum, & Bülthoff, 2005; Schulte-Pelkum et al., 2004) or head-mounted displays (Bakker, Werkhoven, & Passenier, 1999, 2001; Riecke, Schulte-Pelkum, & Bülthoff, 2005). (See, however, Riecke et al., 2002.) Critically, all participants reported the visually simulated turning direction correctly for both trials. Note that this was also the case for the left-right inverters, and can thus be excluded as a potential reason for the left-right inversion. Interestingly, though, when asked to report the direction of the image motion on the projection screen, 2 of 24 and 3 of 24 participants reported the wrong direction in trials 1 and 2, respectively. This suggests that participants paid more attention to the meaning/interpretation of the visually displayed motion than to the motion of the visual pattern on the screen itself. Unexpectedly, a large percentage of the participants were not able to report correctly whether the trajectory was equilateral or non-equilateral: 15 of the 24 participants (62.5%) made errors for trial 1 (equilateral), and 4 of 24 participants (16.7%) responded erroneously for trial 2 (non-equilateral). This suggests that at least some of the participants might not have been able to disambiguate between the equilateral and non-equilateral paths in the main experiment either, which might have contributed to the overall large systematic errors and lack of significant influence of s_1.

3.2.2.4 3D Computer Game Experience. Similar to the first experimental series, fewer left-right inverters had experience with 3D computer games (1/11, as compared to 5/13 for the non-inverters), but this trend reached only marginal significance ($\chi^2(1, N = 24) = 2.74$, $p = .098$). General computer usage did not differ

8. Due to technical problems, the data from seven participants was missing for the self-rated spatial orientation ability and the rated certainty that one's response was correct.

significantly between inverters and non-inverters (cf. Figure 14), corroborating the findings from the first experimental series.

3.2.2.5 Gender Effects. The first experimental series showed a significant gender difference between the group of left-right inverters, which was predominantly female (five out of six inverters were female) and the non-inverters, which were predominately male (11/16 male). The current study revealed a similar trend. Left-right inverters were predominately female (7/11), whereas non-inverters were predominately male (8/13). Note, however, that this trend did not reach significance ($\chi^2(1, N = 24) = 1.51, p = .22$). Male participants were on average 4.5 years older than females (cf. Figure 14), whereas the inverters and non-inverters showed no significant age difference. Despite gender being known to correlate with many spatial abilities (see Coluccia & Louse, 2004; and Lawton & Morrin, 1999, for comprehensive reviews), gender was not associated with lower mental spatial abilities in the current study, neither in terms of the paper-and-pencil questionnaire (Stumpf & Fay, 1983) nor the self-evaluation (cf. Figure 14).

3.2.2.6 Relation between 3D Computer Game Experience, Gender, and Left-Right Inversion? As both gender and 3D computer game experience correlated with the occurrence of left-right inversion in the first experimental series, one might hypothesize that experience with 3D computer games might somehow be able to overcome the gender effect. That is, are female participants with 3D computer game experience less likely to show left-right inversion errors? Clearly, a much larger participant population would be needed to properly address this issue. From the current data, we can only observe that for both the first and second experimental series, there was only one women each who had any 3D computer game experience, and both of those women did not show left-right inversions. Note that even if a larger participant population were to be tested, such correlations would, of course, not necessarily imply any causal relation between 3D computer game experience and a lack of left-right inversion.

3.3 Intermediate Conclusions

In conclusion, we were able to replicate the occurrence of consistent left-right inversion errors for about 40% of naive, psychophysically inexperienced observers in the second experimental series. In addition to producing consistent left-right errors, the inverters also showed overall larger pointing variability and errors that cannot be simply attributed to the left-right errors themselves. Furthermore, left-right errors were associated with lower overall mental spatial abilities in both experimental series. While the first experimental series found significant correlations between left-right inversion and both gender and the lack of 3D computer game experience, these correlations did not reach significance in the second experimental series—although there were trends in the same direction.

One of the most striking results of the second experimental series was the relatively poor overall performance for all but a few participants, especially given that they had advance information about the exact turning angle and were able to perform virtually flawlessly in a real-world pre-test. Furthermore, the VR point-to-origin task was still perceived as surprisingly difficult, which was corroborated by the relatively long response times. In fact, not a single participant reported having any kind of natural or intuitive spatial orientation during the VR experiments, not even the VR-experienced lab members in the encoding control experiment of the first experimental series. Thus, at least for the current VR setup, it seems as if optic flow itself might not be sufficient for enabling natural and intuitive spatial orientation or automatic spatial updating for even the most basic and seemingly trivial trajectories, despite the additional information about the turning angles.

4 General Discussion and Conclusions

4.1 Potential Factors Underlying Left-Right Inversions

In the following, we will discuss a number of potential factors that might have contributed to the observed left-right inversions. Potential systematic left-right errors in the data analysis were excluded by producing consistent dummy data and running the identical analysis. The high consistency and low circular standard deviation observed for many of the participants argues against a high random error as a potential explanation for the observed left-right inversion. Furthermore, the overall low circular standard deviations and the high accuracy observed for the 0° conditions (which does not involve any significant challenges in terms of encoding or mental spatial computation) argues against a noteworthy contribution of execution errors due to the pointing procedure or apparatus. Moreover, participants in the real-world pre-test pointed with high accuracy and minimal systematic error, corroborating the low execution error. One might also argue that left-right inversion could have been caused by a systematic misperception of the direction of the visually simulated turns (e.g., rightward optic flow might not have been perceived as simulating a left turn, but erroneously a right turn). To test this hypothesis, we added two post-experimental trials to the second experimental series where participants were explicitly asked about the direction of the simulated self-rotation. All participants reported, however, the correct direction for both trials, suggesting that the turning direction was correctly perceived throughout the experiment and did not cause the observed left-right inversion. One might also imagine that the left-right inverters were simply irreversibly lost and/or confused. This would predict that the inverters responded rather randomly or at least without much correlation to the turning angle, and should thus show less consistent data than the non-inverters. There was, however, a clear correlation between the experimental manipulation of the turning angle and the inverters' responses (see, e.g., Figure 7). In fact, the inverters' pointing responses closely mimic the non-inverters' responses if left-right mirrored. Furthermore, both the absolute pointing error and the circular standard deviations were well below chance level. In fact, the consistency of participants' pointings was surprisingly high, and not noticeably lower for the two-segment task, which was clearly more complex than the one-segment task and thus more prone to disorientation (see Figures 6 and 13). Moreover, the inverters did not show higher circular standard deviations than the non-inverters for the more complex two-segment task (see Tables 1 and 2). Finally, if participants were truly lost or confused, one would expect inconsistent or random left-right responses. Each of the participants showed, however, consistent pointing directions—despite in the consistently left-right reversed direction for the inverters. A systematic misunderstanding of the experimental instructions for the left-right inverters seems also quite unlikely, given the extensive instructions and in particular the data from the real-world pre-experiment (where they performed virtually flawlessly), the two post-experimental trials, and the extensive debriefing. There are, however, two different hypotheses that might both be able to explain the systematic left-right inversions: failure to update one's heading or simply left-right swap errors.

The consistent left-right inversions observed in both experimental series of the current study bear resemblance to differences in pointing strategies observed by Gramann et al. (2005) using point-to-origin experiments in desktop VR. In their study, only 20 of the 43 participants (accumulated over three experiments) updated their heading according to the visual turns (so-called turners). The other 23 participants responded as if they somehow failed to update the visually displayed heading changes and thus responded as if still facing the original direction (non-turners). These left-right inversions observed in VR studies resemble real-world data from imagined walking experiments using two-segment excursion (Avraamides et al., 2004; Klatzky et al., 1998). Instead of using a pointing method, participants in these real-world studies were asked to turn to face the origin as if they had actually walked the excursion trajectory. Participants failed to update heading changes during the imagined rotation between the first and second segment, and responded as if standing at the to-be-imagined location, but still facing the initial orientation. Similar failures to update rotations were observed by Klatzky and colleagues when only optic flow information presented on an HMD indicated the trajectory or when participants watched another person walk the excursion path. Only when participants actually walked the path or at least physically executed the turn between the two segments did they properly incorporate the rotation, which corroborates the often posited importance of physical motion cues for automatic spatial updating (Farrell & Robertson, 1998; May & Klatzky, 2000; Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004).

Gramann et al. (2005) argued that the turners in their study used an egocentric strategy, whereas the non-turners used an allocentric strategy. Following this line of reasoning, one might be tempted to conclude that the left-right inverters in the current experiment were non-turners and hence did not update their facing direction according to the visual stimulus at all. The non-inverters would correspondingly be categorized as turners. To give an example, if participants are presented with a two-segment isosceles path containing a 60° left turn, the turners would point 150° to the *left*, whereas the non-turners would be expected to point 150° to the *right* (see Figure 4).

Thus, it seems like failure to update heading changes might be a possible explanation for the left-right errors observed in the current study. One participant of the second experimental series did, in fact, realize himself during the debriefing that he completely forgot to update his heading during the VR tests—and was quite surprised by this fact. Three further left-right inverters did not realize themselves that they forgot to update their heading, but when this possibility was mentioned to them at the end of the debriefing they agreed that this is what happened during the VR tests. The other left-right inverters of the second experimental series, however, did not explicitly confirm a failure to update their heading, and there is, indeed, another possible explanation that is consistent with the data from Gramann et al. (2005), Klatzky et al. (1998), and the current study, namely, that the left-right inverters were initially uncertain about the correct response, or somehow puzzled or distracted by the visual simulation, and initially picked the left-right mirrored response and then continued to do this, resulting in consistent left-right swap errors. For equilateral paths ($s_1 = s_2$), both hypotheses produce, in fact, identical prediction and can thus not be directly disambiguated. Unfortunately, the debriefing in the current study showed that participants were not able to reliably determine whether $s_1 = s_2$ or $s_1 > s_2$, such that the data from the $s_1 > s_2$ trials cannot be directly used to disambiguate between the two hypotheses. We are currently conducting experiments with much larger, clearly perceivable differences between s_1 and s_2 to test both hypotheses. There are, however, two pieces of evidence that seem to suggest that, to say the least, not all of the current left-right errors can be fully explained by failures to update one's heading.

4.1.1 Left-Right Inverters Did Update Their Heading for the One-Segment Experiment. First, for the one-segment experiment, simple failure to update one's heading would predict that the left-right inverters should not show any sensitivity to turning angles and instead always point directly backward (180°), which was clearly not the case (see Figures 5 and 7). Instead, most left-right inverters showed a clear sensitivity to turning angles and produced the same consistent left-right errors as for the two-segment experiment. Thus, it seems that the left-right inverters either did not simply fail to update their heading, or they suddenly and consistently switched their response strategy for the one-segment experiment—potentially because they realized that always pointing to 180° is somewhat absurd and might make them look silly. Such a switch in response strategy was, in fact, observed by Avraamides et al. (2004) when participants were asked to indicate the direction to the origin of locomotion after an imagined excursion using a bodily response (turning to face the origin as if they had actually walked the path). For two-segment excursions, participants responded as if they failed to update their heading, just as in the study by Klatzky et al. (1998). When presented with two additional one-segment paths, however, the same participants apparently switched response strategy and correctly incorporated the rotation and turned to face the origin. This is a critical difference from the current study, where the left-right inverters did not switch to a correct strategy for the one-segment experiment (if they switched strategy at all), but instead produced the same left-right pointing errors as before. In subsequent two-segment trials, the same participants of Avraamides et al. showed, however, left-right errors and responded as if failing to update their heading, thus replicating findings of their main experiment and of Klatzky et al. That is, being able to update heading for the one-segment experiment apparently did not transfer to the two-segment trials for the Avraamides et al. study. It should be noted, though, that participants performed only two trials each in the one- and two-segment conditions of the control experiment of Avraamides et al., and that exposure to more trials might eventually have led participants to notice their mistake and correct their responses for the two-segment task.

4.1.2 Some Left-Right Inverters Pointed Forward and Thus Did Update Their Heading. Second, if the left-right inverters simply failed to update their heading, they should always point backward (i.e., egocentric pointing directions should always be in the interval $\alpha \in [90°, 270°]$), and never forward ($\alpha < 90°$ or $\alpha > 270°$), as the origin of locomotion would always be behind them (unless $s_2 > s_1$, which was never the case for the current studies, and participants were aware of that fact). As Figures 7 and 12 clearly show, however, there were several left-right inverters[9] who did, in fact, point consistently forward for the largest turns, and thus produced egocentric pointing angles α outside of the interval [90°, 270°]. This suggests that the left-right inversion observed in the current study cannot be fully and for all participants explained by a simple failure to update rotations that are not physically performed.

4.1.3 Two Subgroups of Left-Right Inverters? So what can we conclude about the origin of the observed consistent left-right inversions in VR? The above discussion suggests that the group of left-right inverters is not homogeneous, but instead clusters into two subgroups.

On the one hand, there are some left-right inverters that apparently did update the visually displayed heading changes, but for some reason produced a left-right mirrored response, despite correctly perceiving the simulated turning direction. Post-experimental debriefing suggests that they might initially have been uncertain about the proper response and for whatever reason picked the wrong, left-right reversed, strategy and later continued to employ that same strategy.

On the other hand, the larger subgroup of left-right inverters produced data that are at least roughly consistent with a failure to update the visually presented heading changes—but only if one accepts that they all somehow switched their response strategy for the one-segment experiment and immediately adopted a new strategy that did incorporate heading changes while keeping the left-right reversals.

4.1.4 Conclusions. No matter what caused the left-right inversions, it seems obvious that updating/computing a homing vector can be quite difficult in VR, potentially because of the lack of physical motion cues, the absence of landmarks, and/or the fact that the pointing target remains mostly outside of the field of view of the visual display. Furthermore, the observed qualitative errors might simply disappear if a verbal response is used instead of a bodily response like pointing, as was the case in the imagined walking study by Avraamides et al. (2004). According to the sensorimotor interference hypothesis, the perceived discrepancy between the physical orientation (which remained unchanged) and the to-be-imagined or visually simulated orientation creates an interference at the response level if (and only if) a bodily response like pointing or turning one's body is used (May, 1996, 2004; Presson & Montello, 1994; Riecke et al., 2007; Riecke & McNamara, 2007; Wang, 2005; Wraga, 2003). Hence, VR users might well be able to update some kind of "cognitive heading" deliberately, which allows for verbal or other responses that are somewhat detached from one's bodily reference frame (Avraamides et al., 2004). Whenever a more natural, embodied response like pointing, grasping, or turning is required, however, VR systems that provide neither physical motion cues nor useful landmarks might be insufficient. This might severely limit the effectiveness and user acceptance of VR at large, in particular for tasks where robust and effortless (i.e., natural) spatial orientation is essential.

9. Namely, participant ID 1 in the first experimental series and participant IDs 4, 5, 8, and 22 in the second experimental series.

4.2 Automatic Spatial Updating Versus Effortful Cognitive Processing

The small overall response times and the lack of an increase in response time for increasing turning angles and number of path segments suggests that the participant already performed some mental spatial computations like updating a homing vector during the simulated motion, and did not wait until the end of the trajectory was reached. The overall large circular standard deviations and the striking qualitative errors (consistent left-right inversions observed for about 40% of the participants) provide, however, strong evidence that spatial updating was by no means automatic and effortless, despite apparently occurring online during the simulated motion. This lack of automatic spatial updating was corroborated by participants' subjective ratings of task difficulty: In the first experimental series, only one participant rated the task as easy, whereas eight participants rated the task as medium difficult and seven as quite difficult. Some participants even mentioned explicitly that they were often unsure which direction (left or right) they had to point to—something that is not observed for comparable real-world tasks, even with eyes closed. The fact that the left-right inverters also showed significantly lower mental spatial abilities further suggests that participants were not able to use automatic spatial updating, but instead had to resort to more cognitive strategies like abstract mental spatial reasoning. Even in the one-segment encoding control experiment, where experienced psychophysical observers were provided with explicit verbal information about the turning angle and direction, the task was rated as "surprisingly extremely difficult."

4.3 Conclusions and Guidelines for VR System Designers

Optic flow has been extensively studied during the last decades, and the literature suggests that optic flow can be used to solve a number of tasks including, for example, heading estimation, estimation of distance and turns, and navigation including path integration (Lappe, Bremmer, & van den Berg, 1999; Riecke et al., 2002; W. H. Warren, Kay, Zosh, Duchon, & Sahuc, 2001; R. Warren & Wertheim, 1990). Even though there typically are some systematic errors in the perception of optic flow, one might be tempted to conclude that just about any task that involves self-motion can, in principle, be performed on the basis of optic flow. This study showed, however, that even an immersive, high-resolution video projection setup is not necessarily sufficient for enabling quick and intuitive spatial orientation and automatic spatial updating when only optic flow cues without any landmarks are available. In the following, we would like to discuss potential underlying reasons and provide some guidelines for VR system designers.

The fact that observers in the current study did not actively control the visually displayed motions can presumably be excluded as a contributing factor, as previous studies demonstrated that actively executing a motion is not required for automatic spatial updating, at least for physical motions (Klatzky et al., 1998; Wang & Simons, 1999; Wraga et al., 2004). It is interesting to note that presence/immersion did not seem to play any significant role in the task used, even though one might imagine that a higher degree of immersion might be able to reduce the interference between the physical and visually simulated orientation and thus facilitate the updating of heading, as proposed by Riecke and von der Heyde (2002) and Riecke and McNamara (2007).

A control study by Klatzky et al. (1998, footnote on p. 297) supports the hypothesis that the interference between the physical and visually simulated orientation might disrupt the updating of heading in VR. In a turn-to-face-origin paradigm where excursion trajectories were visually displayed through an HMD, performance improved when participants were disoriented beforehand and thus less aware of the surrounding physical environment. That is, the failure to update visually presented rotations in VR was less pronounced for disoriented participants. Those and our own data led us to posit that the absence of such interference might be an essential factor or even a necessary prerequisite for enabling automatic spatial updating and natural spatial orientation in VR (at least when physical motion cues are absent and the virtual environment does not provide sufficient landmarks; Riecke & von der Heyde, 2002;

Riecke et al., 2007; Riecke & McNamara, 2007). As automatic spatial updating and natural spatial orientation imply a lack of left-right inversion errors, using a simple experimental paradigm like the point-to-origin paradigm of the current study can provide a simple yet effective benchmarking of a given VR setup and application. When left-right inversions occur, the given VR environment will not be able to afford natural spatial orientation and automatic, path-integration-based spatial updating.

One issue that should be carefully considered when attempting to improve spatial orientation in VR is the type and properties of the visual display device. Even though the current study was not designed to investigate this issue, the literature suggests that not only the visual FOV, but also the display type and geometry are critical factors for improving the effectiveness of VR systems (Arthur, 2000; Riecke, Schulte-Pelkum, & Bülthoff, 2005; Schulte-Pelkum et al., 2004; Tan, Gergle, Scupelli, & Pausch, 2006). Tan and colleagues demonstrated, for example, that using a physically larger display (193 × 145 cm vs. 36 × 26 cm) can enhance performance in a variety of spatial tasks, even though the physical FOV subtended by the two displays was identical (31° × 24°) (Tan et al., 2006). Compared to the larger display of their study, the current display was slightly smaller (168 × 126 cm), but provided a much larger FOV (84° × 63°) and a higher resolution (1,400 × 1,050 pixel compared to 1,024 × 768). A review of all the display factors affecting spatial orientation in VR would go beyond the scope of the current paper; as a rough guideline, though, human spatial orientation in VR typically benefits from a large FOV, and physically large projection screens (in particular if they are curved around the observer) tend to outperform flat or desktop displays and HMDs (Arthur; Darken & Peterson, 2002; Péruch & Gaunet, 1998; Riecke et al., 2002; Riecke, von der Heyde, & Bülthoff, 2005; Riecke, Schulte-Pelkum, & Bülthoff, 2005; Ruddle & Jones, 2001; Schulte-Pelkum et al., 2004; Tan et al., 2006).

A different approach to improve spatial orientation in VR is to train the users on the task. Several studies showed that extensive feedback training can be employed to increase accuracy for a given task (Gramann et al., 2005; Lawton & Morrin, 1999; Riecke et al., 2002; Wiener & Mallot, 2006). The relatively long response times, and the high cognitive demand and rated task difficulty that are often observed in those studies suggest, however, that participants still do not necessarily have a robust and effortless spatial orientation comparable to real-world performance, despite the training. This is in agreement with spatial updating studies that found a lack of automatic spatial updating whenever simulated self-rotations were not physically performed (Chance et al., 1998; Klatzky et al., 1998; Wraga et al., 2004) and/or only optic flow stimuli without useful landmarks were presented (Klatzky et al.; Riecke et al., 2007; Wraga et al.). That is, as a general guideline, VR users should be allowed to physically locomote through the environment or at least physically rotate whenever feasible (Chance et al.; Klatzky et al.; Ruddle & Lessels, 2006; Wraga et al.). Razzaque and colleagues suggested that a limited walking area could still convey much larger virtual spaces if the virtual scene is automatically and imperceptibly rotated such that users are guided away from the physical walls of the lab and/or gain factors between virtual and real rotations are used (Razzaque, Kohn, & Whitton, 2001; Razzaque, 2005). Despite being an interesting approach, this "redirected walking" technique has some practical limitation and has not been systematically tested in terms of spatial orientation, though. Similarly, resetting the observer's virtual position once coming close to the limits of the physical space (e.g., Williams et al., 2007) might work for some applications, but is probably too disruptive for general usage.

If physical walking or turning is unfeasible due to restrictions of lab space, technical effort, or simply money, providing at least *some* qualitatively correct physical motion cues (e.g., by leaning or moving a few centimeters in the direction of intended travel) can still be used as a simple yet effective means to improve spatial orientation and self-motion perception in VR, compared to mere joystick navigation without any physical motion (Peterson, Wells, Furness, & Hunt, 1998; Riecke, 2006). Similarly, providing a small physical motion at the onset of the visually simulated *passive* motion has been shown to improve self-motion illusions for both rotations and translations (Riecke et al., 2006; Wong & Frost, 1981).

Even if no physical motion cues are provided at all, visual cues alone *can,* under certain conditions, still be sufficient for enabling natural, real-world-like spatial orientation and automatic spatial updating if a naturalistic, user-known visual stimulus that includes useful landmarks is used (Riecke, von der Heyde, & Bülthoff, 2005; Riecke et al., 2007). In general, however, it remains an open question and a considerable technological and scientific challenge to determine whether visual information devoid of any landmarks might in principle be capable of enabling robust and effortless, natural spatial orientation in simulated environments. If so, this would be of substantial impact both for our understanding of human spatial orientation and for the design of human-centered, lean, yet effective, self-motion simulators.

We propose that using a simple-to-implement psychophysical test such as the rapid point-to-origin paradigm of the current study and quantifying the occurrence of qualitative errors like left-right inversions might be a powerful way of assessing the perceptual and behavioral effectiveness of a given VR system, and, in particular, its capability in enabling natural and unencumbered spatial behavior and performance. Such straightforward experiments could, in fact, be used as a simple benchmark for evaluating different self-motion simulations and VR setups from a human-centered perspective.

Acknowledgments

This research was funded by Vanderbilt University, the Max Planck Society, and the EU grant POEMS-IST-2001-39223 (see www.poems-project.info). I would like to thank Jan Malte Wiener for great discussions and programming the experiment. Also, I appreciate the assistance of Jan Oliver Him and Naima Laharnar in running the experiment.

References

Arthur, K. W. (2000). *Effects of field of view on performance with head-mounted displays.* Unpublished doctoral dissertation, Department of Computer Science, University of North Carolina, Chapel Hill. Available at ftp://ftp.cs.unc.edu/pub/ publications/techreports/00-019.pdf. Accessed November, 2007.

Avraamides, M. N., Klatzky, R. L., Loomis, J. M., & Golledge, R. G. (2004). Use of cognitive versus perceptual heading during imagined locomotion depends on the response mode. *Psychological Science, 15*(6), 403–408.

Bakker, N. H., Werkhoven, P. J., & Passenier, P. O. (1999). The effects of proprioceptive and visual feedback on geographical orientation in virtual environments. *Presence: Teleoperators and Virtual Environments, 8*(1), 36–53.

Bakker, N. H., Werkhoven, P. J., & Passenier, P. O. (2001). Calibrating visual path integration in VEs. *Presence: Teleoperators and Virtual Environments, 10*(2), 216–224.

Batschelet, E. (1981). *Circular statistics in biology.* San Diego, CA: Academic Press.

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*(2), 168–178.

Coluccia, E., & Louse, G. (2004). Gender differences in spatial orientation: A review. *Journal of Environmental Psychology, 24*(3), 329–340.

Darken, R. P., & Peterson, B. (2002). Spatial orientation, wayfinding and representation. In K. M. Stanney (Ed.), *Handbook of virtual environments* (pp. 493–518). Mahwah, NJ: Lawrence Erlbaum.

Darken, R. P., & Sibert, J. L. (1996). Navigating large virtual spaces. *International Journal of Human-Computer Interaction, 8*(1), 49–71.

Farrell, M. J., & Robertson, I. H. (1998). Mental rotation and the automatic updating of body-centered spatial relationships. *Journal of Experimental Psychology—Learning Memory and Cognition, 24*(1), 227–233.

Gramann, K., Muller, H. J., Eick, E. M., & Schonebeck, B. (2005). Evidence of separable spatial representations in a virtual navigation task. *Journal of Experimental Psychology—Human Perception and Performance, 31*(6), 1199–1223.

Grant, S. C., & Magee, L. E. (1998). Contributions of proprioception to navigation in virtual environments. *Human Factors, 40*(3), 489–497.

IJsselsteijn, W. A. (2004). *Presence in depth.* Unpublished doctoral dissertation, Technische Universiteit Eindhoven, The Netherlands.

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*(3), 349–374.

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*(4), 293–298.

Klatzky, R. L., Loomis, J. M., & Golledge, R. G. (1997). Encoding spatial representations through nonvisually guided locomotion: Test of human path integration. In D. Medin (Ed.), *The psychology of learning and motivation,* Vol. 37 (pp. 41–84). San Diego, CA: Academic Press.

Lappe, M., Bremmer, F., & van den Berg, A. V. (1999). Perception of self-motion from visual flow. *Trends in Cognitive Sciences, 3*(9), 329–336.

Lawson, B. D., Graeber, D. A., Mead, A. M., & Muth, E. R. (2002). Signs and symptoms of human syndromes associated with synthetic experiences. In K. M. Stanney (Ed.), *Handbook of virtual environments* (pp. 589–618). Mahwah, NJ: Lawrence Erlbaum.

Lawton, C. A., & Morrin, K. A. (1999). Gender differences in pointing accuracy in computer-simulated 3D mazes. *Sex roles, 40*(1–2), 73–92.

Loomis, J. M., Da Silva, J. A., Philbeck, J. W., & Fukusima, S. S. (1996). Visual perception of location and distance. *Current Directions in Psychological Science, 5*(3), 72–77.

Loomis, J. M., Klatzky, R. L., Golledge, R. G., & Philbeck, J. W. (1999). Human navigation by path integration. In R. G. Golledge (Ed.), *Wayfinding behavior: Cognitive mapping and other spatial processes* (pp. 125–151). Baltimore: Johns Hopkins.

Marlinsky, V. V. (1999). Vestibular and vestibulo-proprioceptive perception of motion in the horizontal plane in blindfolded man—III. Route inference. *Neuroscience, 90*(2), 403–411.

May, M. (1996). Cognitive and embodied modes of spatial imagery. *Psychologische Beiträge, 38*(3–4), 418–434.

May, M. (2004). Imaginal perspective switches in remembered environments: Transformation versus interference accounts. *Cognitive Psychology, 48*(2), 163–206.

May, M., & Klatzky, R. L. (2000). Path integration while ignoring irrelevant movement. *Journal of Experimental Psychology—Learning, Memory and Cognition, 26*(1), 169–186.

Muehl, K. A., & Sholl, M. J. (2004). The acquisition of vector knowledge and its relation to self-rated direction sense. *Journal of Experimental Psychology— Learning, Memory and Cognition, 30*(1), 129–141.

Péruch, P., & Gaunet, F. (1998). Virtual environments as a promising tool for investigating human spatial cognition. *Current Psychology of Cognition, 17*(4–5), 881–899.

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

Peterson, B., Wells, M., Furness, T., III, & Hunt, E. (1998). The effects of the interface on navigation in virtual environments. *Proceedings of Human Factors and Ergonomics Society 1998 Annual Meeting.*

Presson, C. C., & Montello, D. R. (1994). Updating after rotational and translational body movements: Coordinate structure of perspective space. *Perception, 23*(12), 1447–1455.

Prothero, J. D. (1998). *The role of rest frames in vection, presence and motion sickness.* Unpublished doctoral dissertation, University of Washington, Seattle. Available at www.hitl.washington.edu/ publications/r-98-11/. Accessed November 2007.

Razzaque, S. (2005). *Redirected walking.* Unpublished doctoral dissertation, University of North Carolina at Chapel Hill.

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

Riecke, B. E. (2006). Simple user-generated motion cueing can enhance self-motion perception (vection) in virtual reality. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST)*, 104–107.

Riecke, B. E., Cunningham, D. W., & Bülthoff, H. H. (2007). Spatial updating in virtual reality: The sufficiency of visual information. *Psychological Research, 71*(3), 298–313.

Riecke, B. E., & McNamara, T. P. (2007). An integrative theory of spatial orientation in the immediate environment. *Proceedings of the 29th Conference of the Cognitive Science Society (CogSci).*

Riecke, B. E., Schulte-Pelkum, J., & Bülthoff, H. H. (2005). Perceiving simulated ego-motions in virtual reality—Comparing large screen displays with HMDs. *Proceedings of the SPIE,* 5666, 344–355.

Riecke, B. E., Schulte-Pelkum, J., & Caniard, F. (2006). Visually induced linear vection is enhanced by small physical accelerations. *Proceedings of the 7th International Multisensory Research Forum (IMRF).* Available at www.kyb.mpg.de/ publication.html?publ=3901.

Riecke, B. E., Schulte-Pelkum, J., Caniard, F., & Bülthoff, H. H. (2005). Towards lean and elegant self-motion simulation in virtual reality. *Proceedings of IEEE Virtual Reality 2005,* 131–138.

Riecke, B. E., van Veen, H. A. H. C., & Bülthoff, 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.

Riecke, B. E., & von der Heyde, M. (2002). *Qualitative modeling of spatial orientation processes using logical propositions: Interconnecting spatial presence, spatial updating, piloting, and spatial cognition* (Tech. Rep. No. 100). Max Planck Institute for Biological Cybernetics, Tuebingen, Germany. Available at www.kyb.mpg.de/publication.html?publ=2021.

Riecke, B. E., von der Heyde, M., & Bülthoff, H. H. (2005). Visual cues can be sufficient for triggering automatic, reflex-like spatial updating. *ACM Transactions on Applied Perception (TAP), 2*(3), 183–215.

Riecke, B. E., & Wiener, J. M. (2007). Can people not tell left from right in VR? Point-to-origin studies revealed qualitative errors in visual path integration. *Proceedings of IEEE Virtual Reality 2007,* 3–10.

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., & Jones, D. M. (2001). Movement in cluttered virtual environments. *Presence: Teleoperators and Virtual Environments, 10*(5), 511–524.

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., Payne, S. J., & Jones, D. M. (1998). Navigating large-scale "desk-top" virtual buildings: Effects of orientation aids and familiarity. *Presence: Teleoperators and Virtual Environments, 7*(2), 179–192.

Sadalla, E. K., & Montello, D. R. (1989). Remembering changes in direction. *Environmental Behavior, 21*(3), 346–363.

Sadowski, W., & Stanney, K. (2002). Presence in virtual environments. In K. M. Stanney (Ed.), *Handbook of virtual environments* (pp. 791–806). Mahwah, NJ: Lawrence Erlbaum.

Schulte-Pelkum, J., Riecke, B. E., von der Heyde, M., & Bülthoff, H. H. (2004). *Influence of display device and screen curvature on perceiving and controlling simulated egorotations from optic flow* (Tech. Rep. No. 122). Max Planck Institute for Biological Cybernetics, Tuebingen, Germany. Available at www.kyb.mpg.de/publication.html?publ=2574.

Shepard, R. N., & Metzler, J. (1971). Mental rotation of 3-dimensional objects. *Science, 171*(3972), 701–703.

Sholl, M. J. (1989). The relation between horizontality and rod-and-frame and vestibular navigational performance. *Journal of Experimental Psychology—Learning, Memory and Cognition, 15*(1), 110–125.

Stumpf, H., & Fay, E. (1983). *Schlauchfiguren—Ein Test zur Beurteilung des räumlichen Vorstellungsvermögens.* Göttingen, Germany: Hogrefe.

Tan, D. S., Gergle, D., Scupelli, P., & Pausch, R. (2006). Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction, 13*(1), 71–99.

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.

Wang, R. F. (2005). Beyond imagination: Perspective change problems revisited. *Psicologica, 26*(1), 25–38.

Wang, R. F., & Simons, D. J. (1999). Active and passive scene recognition across views. *Cognition, 70*(2), 191–210.

Warren, R., & Wertheim, A. H. (Eds.). (1990). *Perception and control of self-motion.* Mahwah, NJ: Lawrence Erlbaum.

Warren, W. H., 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.

Wiener, J. M., & Mallot, H. A. (2006). Path complexity does not impair visual path integration. *Spatial Cognition and Computation, 6*(4), 333–346.

Williams, B., Narasimham, G., Rump, B., McNamara, T. P., Carr, T. H., Rieser, J., et al. (2007). Exploring large virtual environments with an HMD when physical space is limited. In *ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization.*

Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments, 7*(3), 225–240.

Wong, S. C. P., & Frost, B. J. (1981). The effect of visual-vestibular conflict on the latency of steady-state visually induced subjective rotation. *Perception & Psychophysics, 30*(3), 228–236.

Wraga, M. (2003). Thinking outside the body: An advantage for spatial updating during imagined versus physical self-rotation. *Journal of Experimental Psychology—Learning, Memory and Cognition, 29*(5), 993–1005.

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.