The Influence of Visual Structure and Physical Motion Cues on Spatial Orientation in a Virtual Reality Point-to-Origin Task

by

Salvar Sigurdarson

B.Sc., University of Iceland, 2007

Thesis Submitted In Partial Fulfillment of the
Requirements for the Degree of
Master of Science

in the
School of Interactive Arts & Technology
Faculty of Communication, Art and Technology

SIMON FRASER UNIVERSITY

Summer 2014

Approval

Name: **Salvar Sigurdarson**

Degree: **Master of Science (Interactive Arts & Technology)**

Title of Thesis: ***The Influence of Visual Structure and Physical Motion Cues on Spatial Orientation in a Virtual Reality Point-to-Origin Task***

Examining Committee: **Chair:** Dr. Philippe Pasquier
Associate Professor

Dr. Bernhard Riecke
Senior Supervisor
Associate Professor

Dr. Alissa Antle
Committee Member
Associate Professor

Dr. Roberta Klatzky
External Examiner
Carnegie Mellon University

Date Defended/Approved:

Ethics Statement

The author, whose name appears on the title page of this work, has obtained, for the research described in this work, either:

a. human research ethics approval from the Simon Fraser University Office of Research Ethics,

or

b. advance approval of the animal care protocol from the University Animal Care Committee of Simon Fraser University;

or has conducted the research

c. as a co-investigator, collaborator or research assistant in a research project approved in advance,

or

d. as a member of a course approved in advance for minimal risk human research, by the Office of Research Ethics.

A copy of the approval letter has been filed at the Theses Office of the University Library at the time of submission of this thesis or project.

The original application for approval and letter of approval are filed with the relevant offices. Inquiries may be directed to those authorities.

Simon Fraser University Library
Burnaby, British Columbia, Canada

update Spring 2010

Abstract

Virtual reality simulators have a serious flaw: Users tend to get lost and disoriented as they navigate. The prevailing opinion is that this is due to the lack of physical motion cues, but a growing body of research challenges this notion. In two experiments, 48 participants estimated their position after passive motions in a virtual environment without landmarks (ranging from pure optic flow to a structured city), by pointing towards the origin of the simulated movement. In half of the trials the visually displayed turns were accompanied by a matching physical rotation. Results showed that while physical rotation cues did not improve spatial orientation performance, structured visuals did. Furthermore, we observed that visuals experienced first by a participant significantly affected spatial orientation performance in subsequent environments. Our findings lend support to the notion that spatial orientation ability in VR may not require physical motion cues, but can be facilitated by a naturalistic and structured environment. This knowledge improves our understanding of how different modalities affect human spatial cognition, and can guide the design of safer and more affordable VR simulators.

Keywords: spatial updating; spatial orientation; virtual reality; point-to-origin; Turner; Non-Turner; egocentric

Acknowledgements

First and foremost, I would like to thank my supervisor Dr. Bernhard E. Riecke, and thesis committee member Dr. Alissa Antle, for their academic input and guidance in writing this thesis. I will also extend heartfelt gratitude to my fellow iSpace collaborators, especially Lonnie Dower and Andrew Milne, for their support, insight and shared suffering on our arduous path towards graduation.

Table of Contents

List of Tables

List of Figures

1. Executive summary

Before we begin I would like to ask you to close your eyes, rotate your body 90 degrees to the left, and point to the nearest door.

How did you do? You probably pointed quite accurately with relatively little effort. But how did you know where to point? You certainly weren't thinking consciously about the door before performing this task, and you could not see it when you pointed. Clearly some part of your mind was keeping track of the door's location, relative to your position, while you rotated with your eyes closed.

This important but often overlooked process that enables us to interact with the world without getting lost or disoriented is called *spatial updating*. It allows us to perceive the world as relatively stable while we are moving, to track the movement of other mobile objects, and to use stable objects as landmarks in our navigation. Spatial updating is not unique to humans; it is a ubiquitous process identified in almost all moving organisms including insects, birds, rodents, and primates.

Although implementation can differ, the spatial updating process functions similarly across species: It constructs and updates a mental representation of the environment using available sensory information. For humans, this includes information from our eyes, ears, muscles and the vestibular system located in the inner ear. The spatial updating process is very robust, and can operate even without vision, as anyone who has navigated to the restroom in the middle of the night can confirm.

Continuously keeping track of surrounding objects as we – and they – move is a computationally complex task. Yet, spatial updating happens automatically and effortlessly. In fact, spatial updating is so automatic that suppressing it takes significant mental effort. You can try this in a similar way to the first test: Close your eyes and point to the nearest door as if you were facing the *opposite* direction. If you are like most people, this task was not quite as effortless as the first one. This is because rather than using the spatial updating process, you now had to actively ignore it.

When we move around in the real world, spatial updating is robust, mostly automatic and relatively effortless. During navigation in a virtual environment however, the process does not trigger as easily, resulting in disorientation and confusion. This is a serious flaw in Virtual Reality (VR), and prevents the technology from fulfilling its potential in uses like simulation, training, design and entertainment. After all, a VR flight simulator that leaves users lost after a few minutes will not do a very good job of simulating real flight.

Why does this disparity exist? What is missing in a virtual environment that is present in a real one? For one, the physical motion cues we receive when we move around in the real world are often absent in VR, as they are too complex and expensive to be fully recreated. Without prohibitively expensive setups, VR users will not physically feel movements in the same way as in the real world, although they might see and hear them in a similar way. This could very well be the answer, and it is indeed the prevailing notion that for spatial updating to occur, physical motion cues are absolutely necessary. However, others have challenged this notion and found that under certain conditions, spatial updating can be triggered using only visual cues, e.g. during rotational movements with a structured visual environment (see subsection 2.3.1 for a detailed discussion).

In this study, we extend this line of research. Using state-of-the-art VR technology we set up experiments that allowed us to further investigate how different types of cues affect the spatial orientation of people as they move through a virtual environment. By measuring participants' spatial orientation after moving along curved paths in virtual environments of varying fidelity, with or without physical rotations to match rotations in the virtual environment, we found that naturalistic and structured visuals may under certain conditions reduce or even eliminate the need for physical rotation cues. In some instances, even a minimal amount of structure in the visual stimulus – such as a thin white line providing advance information on turn direction – can achieve this effect.

Furthermore, our results suggest that when moving in a virtual environment, our choice of reference frame is partly dependent on the visual information initially available to us. Specifically, we found that participants who experienced a low-fidelity environment in their first trial were less likely to spatially orient themselves correctly, and more likely to find physical motion cues helpful, than participants who started in a high-fidelity environment.

The ultimate goal of this research is to enable spatial orientation in virtual environments that is as effective as in the real world. Systematically investigating the conditions under which automatic spatial updating occurs deepens our understanding of human spatial cognition and guides the design of more effective Virtual Reality simulations. The experiments and findings detailed in this study are a step towards that goal.

1.1. Research goal

Our goal with this study is to examine the role of physical motion cues and visual structure on the occurrence of automatic spatial updating in virtual environments. Using a custom-built motion simulator and top-of-the-line VR equipment, we ran two experiments to investigate how different inputs affect our sense of orientation in virtual environments. We address the following research questions:

1. When moving along a path in a virtual city, are physical rotations required to enable automatic spatial updating, or can naturalistic and structured visuals suffice?

2. Does a bare minimum of salient visual features, or structure, suffice to enable automatic spatial updating during visual-only motion in a virtual environment?

2. Background and motivation

2.1. Virtual Reality

2.1.1. What is virtual reality?

The term “Virtual Reality” is familiar to most people, but it is challenging to define precisely. Virtual Reality involves some kind of simulation of the real world, but when does something cease to be VR and start being a simple display? Should we draw the boundaries of VR based on the technology used to create the simulation or the way the user perceives the environment?

Traditionally, Virtual Reality was defined from the viewpoint of the technological hardware constructing it. Ivan Sutherland, widely regarded as one of the fathers of VR, described in a paper the "ultimate display": a "looking-glass into the mathematical wonderland constructed in computer memory" (Sutherland, 1965, p. 1). This vision would later grow to become VR, and the focus on technology would follow. For example, Greenbaum (1992, p. 58) refers heavily to hardware in his definition:

> “Virtual Reality is an alternate world filled with computer-generated images that respond to human movements. These simulated environments are usually visited with the aid of an expensive data suit which features stereophonic video goggles and fiber-optic data gloves.”

Coates (1992) provides a similar definition: "Virtual Reality is electronic simulations of environments experienced via head mounted eye goggles and wired clothing enabling the end user to interact in realistic three-dimensional situations."

Although defining Virtual Reality from a mostly technical standpoint can be very useful to hardware designers, it is less helpful when it comes to describing the experiences and feelings of the human being witnessing the simulated reality. To address the need for a more theoretically useful concept, Steuer (1992) described VR from an experiential standpoint. Building on the concept of "presence" (Gibson, 1979) – which refers to the sense of being in an environment – Steuer defined the term “telepresence” (initially coined by Minsky (1980) in reference to systems that allow for remote manipulation of physical objects) as “the experience of presence in an environment by means of a communication medium.” VR could then be defined without reference to any specific hardware system as “a real or simulated environment in which a perceiver experiences telepresence” (Steuer, 1992).

In this study we focus on the mind experiencing a virtual environment, rather than the technology delivering it. We therefore find Steuer’s definition more appropriate for our purposes, although our findings may also have implications for practical design of VR technology.

2.1.2. What are Virtual Reality’s real-world applications?

Some readers will remember virtual reality from hyped magazine articles and failed promises of the 1990s, and might assume that the idea itself turned out to be a failure. This is not the case. Lawnmower Man aside, VR has been researched quite extensively and is currently used in a multitude of situations where using a real environment is impractical.

Training: The freedom to make mistakes is a crucial part of training, but how do you gain experience if the slightest error can result in tragedy? This is a problem that medical doctors, pilots and many other professions have faced, and one that VR can help solve. By using a simulated environment that is as realistic as possible, trainees have the freedom to make mistakes without real-world consequences, yet they learn skills that ideally transfer to the real environment. The most prominent examples of this use are in flight simulation and surgical training (Aggarwal et al., 2007; Lawson & Riecke, 2014; Ota, Loftin, Saito, Lea, & Keller, 1995; Psotka, 1995).

Education: Outside the realm of professional training, the ability to create a safe simulation of dangerous situations has been harnessed for other educational purposes. For example, immersive game-based VR has been successfully used to help children learn about fire hazards (Smith & Ericson, 2009) and geography (Virvou, Manos, Katsionis, & Tourtoglou, 2002).

Design: Visualizing early prototypes is an important part of the iterative process of design but this is a challenge for large-scale projects, such as the layout of a skyscraper. Models are a common solution to this problem, but recently architects have also started using Virtual Reality simulations. While a miniature model of a building can give an architect valuable insight, an immersive virtual environment affords her experience that the model cannot, such as viewing the building from the inside and observing how sunshine enters through windows and reflects off walls (Frost & Warren, 2000).

Therapy: Treating mental disorders such as phobias often involves exposing patients to their fears in a safe and controlled environment. In some cases it is either infeasible or downright impossible to do this in a real environment, e.g. trauma treatment for soldiers, while maintaining adequate safety and/or control. Enter VR: A simulation of a fear-enhancing situation can trigger real psychological symptoms such as sweating or nausea while affording full control and posing no safety threats. This approach to exposure therapy has been used for phobias such as fear of flying (Mühlberger, Herrmann, Wiedemann, Ellgring, & Pauli, 2001) and fear of heights (Emmelkamp et al., 2002; Rothbaum et al., 1995), but research has also delved into using VR to treat other disorders including sexual dysfunctions and eating disorders (Eichenberg & Wolters, 2012; Scozzari & Gamberini, 2011).

Entertainment: Last but not least, VR has been successfully used to create fantastical environments that excite, thrill and entertain without being dangerous or prohibitively expensive. Examples of this use for VR can been seen in theme parks around the world (Pausch, Snoddy, & Taylor, 1996), and immersive video game systems such as the Oculus Rift (Parkin, 2013).

2.1.3. What are its benefits for use in research?

These benefits of VR also make it a very interesting tool for spatial cognition researchers, especially when it comes to experiments on orientation and navigation, as it allows us to create environments that are difficult, expensive and even completely impossible to create in the real world. We can collect data in real time, as it is generated by participants, and even analyze it on the fly to allow for immediate feedback and adjustments. With VR, experimental design can take a dynamic form, and change even during the experiment, for example based on participant performance.

Last but not least, VR is designed to simulate naturalistic environments. Within these environments, researchers can therefore create experimental conditions that are well defined and can easily be reproduced. This is usually not the case with the real world, where there are multiple external factors out of our control. For example, the weather conditions (clouds, visibility, sunlight), location, objects or sound sources may move around, participants may have previous knowledge of the environment, and the ideal environment might not even exist in the researcher's area. The development of VR has provided the opportunity to tackle these issues.

VR has become a prominent tool in the spatial cognition researcher's toolbox, and for good reason (Loomis, Blascovich, & Beall, 1999; Péruch & Gaunet, 1998).

2.1.4. What are its drawbacks for use in research?

This technology is not a silver bullet though, as VR has inherent flaws that limit its usefulness in many situations. First, the technological setups themselves are still quite expensive and complex in operation, and equipment such as head-mounted displays can be uncomfortable to wear for long times. Experiencing VR for a sustained amount of time can also trigger motion sickness in some people, to the point where they cannot complete the experiment due to discomfort (Howarth, Sharples, Cobb, Moody, & Wilson, 2008).

Finally, VR simply doesn't 'feel' quite *real* yet, especially during navigation in a virtual environment (Péruch & Gaunet, 1998). There, users get disoriented quickly, which can severely diminish the ecological validity of the technology as a research and training tool. A driving simulator that leaves users lost after a few turns might not do a good job of training people to drive in the real world.

Why does this happen in VR but not in the real world? What is missing? In order to address this question, we will first introduce a cognitive process that is essential to all mobile species: *Spatial updating*.

2.2. Spatial updating

2.2.1. What is spatial updating?

As we move through the world, our spatial relationship with the environment changes constantly. Surrounding objects can quickly go from being in front of us to behind us, and locations that were near can become distant. Without some process to easily keep track of these changes, it would be very difficult for us to determine whether the world itself is changing, or merely our position and orientation within it. This would make it almost impossible to navigate the environment in order to find food, shelter and safety; something all animals must do in order to survive.

Fortunately such a cognitive process exists. It is called *spatial updating* (Farrell & Robertson, 1998; Klatzky, Loomis, Beall, Chance, & Golledge, 1998; Presson & Montello, 1994; Rieser, Guth, & Hill, 1982), and plays a central role in our interaction with the world. It allows us to perceive the world as stable and to detect changes in position or orientation of mobile objects.

Spatial updating is not unique to humans. In fact, the ability to keep track of position and orientation during movement has been found in almost all animals tested, including desert ants (Müller & Wehner, 1994) spiders (Mittelstaedt, 1985), hamsters (Etienne, Maurer, & Saucy, 1988), honeybees (Esch & Burns, 1996), geese (Mittelstaedt, 1982), gerbils (Mittelstaedt & Glasauer, 1991) and primates (Skolnick, Ackerman, Hofer, & Weiner, 1980).

Spatial updating is at the very core of our behaviour, and without it we would quite literally be lost.

2.2.2. How does spatial updating work?

Although the implementation may differ between species depending on available sensory information and processing capabilities, the basic mechanism of spatial updating remains the same: It creates and updates an internal representation of space and self-to-object relationships in the organism's brain, using sensory information as inputs. Analyzing data from animals, the literature has identified two distinctive but complementary sub-processes that contribute to navigation: Path integration and piloting.

Path integration, also known as 'dead reckoning' (Darwin, 1873), is the most basic form of spatial updating. It is a process whereby the organism continuously integrates its own acceleration and velocity to calculate the location of an origin during movement. For many species, like desert ants and other insects, this is the predominant type of spatial updating available; as more advanced methods require richer sensory information and more complex processing.

Nevertheless, path integration can be quite precise, as evidenced by the ability of desert ants to travel away from home in twisting and turning paths, but to return home in a relatively straight line (Müller & Wehner, 1994). Path integration is subject to accumulative errors however, as each estimate of position and orientation is relative to the previous one (Wang & Spelke, 2002).

Piloting, also called 'landmark-based navigation', is a high-level form of spatial updating where an organism can use direct sensory information and landmarks (i.e. unique and distinctive objects that provide spatial information) to determine their location and orientation. This can give a much more accurate estimate of one's own position, but also requires richer sensory information, more complex processing and an environment that contains salient landmarks. Piloting has been primarily studied in humans, but studies have shown that non-human animals, including insects, are capable of updating their position by way of this process (Wehner & Müller, 2010).

Although these two processes may be distinct, they are quite complementary. Many animals including human beings can use either, or both, depending on the situation. For example, when you walk around your apartment in darkness you are using path integration – with help from your memory of the apartment layout – to navigate around obstacles like tables, chairs and walls. Even without vision, you know more or less how far you have walked based on cues from your proprioceptive and vestibular systems. However once the lights are turned on, piloting kicks in and complements the path integration process by identifying landmarks that can help correct estimation errors. For example, you might realize that you are closer to the refrigerator than you thought.

2.2.3. Can we control spatial updating?

Spatial updating, whether performed via path integration, piloting or a combination of both, involves complex calculations performed in real time. Furthermore, the task of orienting oneself spatially is frequent and essential for most animals to survive, and is often accompanied by other complex tasks – like hunting, flying or fighting – that would be nigh impossible to perform if cognitive resources were largely spent on staying spatially oriented. Fortunately, spatial updating has been found to be automatic on three levels:

Spatial updating is spontaneous in the sense that it occurs without any intention or instruction. An animal need not be told to perform spatial updating as it moves; it simply happens. This applies to humans as well, as we are usually unaware of the fact that we are keeping track of our relationship with surrounding objects almost all the time.

Spatial updating is obligatory in the sense that it occurs involuntarily and is difficult to suppress (Farrell & Robertson, 1998; May & Klatzky, 2000; Riecke, Cunningham, & Bülthoff, 2007). You can try this by closing your eyes, rotating to face the opposite direction, but then pointing to the nearest door as if you had not rotated at all. You will likely find this much more difficult than the previous demonstration, since you first had to “undo” the spatial updating you performed while rotating.

Spatial updating is effortless in the sense that it is quick, easy and requires little attentional focus or other cognitive resources (Farrell & Robertson, 1998; Rieser, 1989; Wraga, Creem-Regehr, & Proffitt, 2004). It is not resource-free however, as research has shown that manipulation of attentional focus can affect the process (Wang & Brockmole, 2003).

2.2.4. How is spatial updating represented in the brain?

Spatial updating may allow us to stay aware of where we are while moving around, but how is that knowledge represented in the brain? What is the frame of reference used to keep track of surrounding objects and locations? These are questions that have led to the two current models of mental spatial representations: *egocentric* and *allocentric*.

In an *egocentric model*, the reference frame is centered on the observer. Locations of surrounding objects are represented with respect to the particular perspective of a perceiver (Klatzky, 1998), and must therefore be updated whenever she moves or rotates. Positions and orientations of objects in the environment are represented as vectors pointing from the observer to each object, and must thus be updated whenever the observer moves or rotates (Wang & Simons, 1999; Wang & Spelke, 2002).

In the *allocentric model* on the other hand, the reference frame is not centered on any specific observer or location but rather all of them at once. An allocentric reference frame locates points within a framework external to the holder of the representation and independent of her position (Klatzky, 1998), similar to a map. Because of this lack of a single center of the reference frame, movement of the observer requires no different calculation than for any other object, as the observer is 'just another item on the map' (Klatzky, 1998; McNamara, Sluzenski, Rump, & Byrne, 2008).

Although the egocentric and allocentric representations were originally envisioned as competing theories, research now points to a two-system model in which egocentric representations exist in parallel to (rather than instead of) allocentric ones (Avraamides & Kelly, 2008; Burgess, 2006). I discuss this in more detail in subsection 2.3.2 below.

2.2.5. *How can we measure spatial updating?*

As spatial updating is triggered spontaneously, obligatorily and operates effortlessly, it can also be very difficult to measure directly. How does one detect a mental process that the owner might not even be aware of? Currently there are three available approaches to this problem:

Neurophysiological study, where brain imaging is used to investigate the neurophysiological changes that happen in the brain during motion, whether real, simulated or imagined. Implanted sensors have been quite useful for the study of spatial orientation in rats, for example, where specific regions of the brain have been identified as “place cells” and “head direction cells” – cells that fire when the animal as it a specific place or orientation, respectively (Gramann, 2013; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Muller, 1996).

However, there is very limited opportunity to use these invasive methods on humans, and the available resolution of non-invasive brain imaging technologies (like electroencephalography) is still too coarse for effectively studying spatial orientation in humans, although there are some promising studies in the literature suggesting that place cells also exist in the human brain (Ekstrom et al., 2003; Gramann, Müller, Schönebeck, & Debus, 2006).

Behavioural study, where a participant is asked to perform a spatial task after translations and/or rotations, and the response is recorded. An established version of this approach is having participants indicate an estimate of their own position relative to an unseen landmark or location, either by pointing physically or describing it verbally. The former is commonly called a rapid pointing paradigm, and has been used effectively in many studies (Klatzky et al., 1998; Riecke & Bülthoff, 2004; Riecke, Heyde, & Bülthoff, 2005; Riecke, 2003; Wan, Wang, & Crowell, 2009; Wang, 2004; Wartenberg, May, & Péruch, 1998; Wiener & Mallot, 2006).

The line of reasoning for this goes: Only if the participant’s mental spatial representation was already automatically updated when she arrived at the new position or orientation, can she give quick, intuitive and accurate spatial answers. This applies whether the movement was real, simulated (e.g. virtual reality) or imagined.

The type of spatial task performed in spatial updating studies varies, but most often they are asked to indicate an estimate of an unseen object's location relative to their own actual, simulated or imagined position and orientation. This can be the starting point of the movement (i.e. 'home'), or some other item that the participant could see before and/or during movement, but not at the new position/orientation. The indication can be verbal ("2 o'clock", "100 degrees left"), translational (travel back to the origin), or physical (point to target).

In the experiments described in this study, we had participants move to a new position and orientation (in a virtual environment), and then asked them to point (with their hand) back to the original location ("home" or "origin").

Introspective study, where a participant is asked to describe feelings and thoughts after performing spatial trials of some kind. As an introspective approach, this method can give insight into spatial updating that quantitative methods cannot. Measuring pointing errors and response times is useful, but simply asking participants how they felt about different stimuli or what strategies they used during trials can provide a valuable supplement to quantitative data.

In experiments described in this study, we used additional introspective measures to collect thoughts and feelings of participants after completing all tasks. In this way, quantitative and qualitative methods are mixed in order to extract as much useful information as possible and analyze it in context (Johnson & Onwuegbuzie, 2004).

2.3. Spatial updating in virtual reality

At this point we know that spatial updating is a robust and mostly automatic process that allows humans and other organisms to navigate their environments without getting lost. When it comes to virtual environments however, users tend to get lost and disoriented quickly (Grant & Magee, 1998; Péruch & Gaunet, 1998). It seems that spatial updating does not get triggered as automatically in VR as it does in the real world. Why? What is missing in VR? In this section, I will review the significant body of literature that discusses this topic.

2.3.1. *Are physical motion cues necessary for effective spatial updating?*

Prevailing notion: Yes

One of the seminal studies on this topic was conducted by Klatzky et al. (1998), where the researchers investigated how participants updated their mental representation of space during real, imagined and simulated movement. Participants were exposed to a two-segment path with a turn between segments (10, 50, 90, 130 or 170 degrees in either direction), and responded by turning to face the origin as they would if they had walked the path and were at the end of the second segment. See Figure 1 for an illustration of this task.

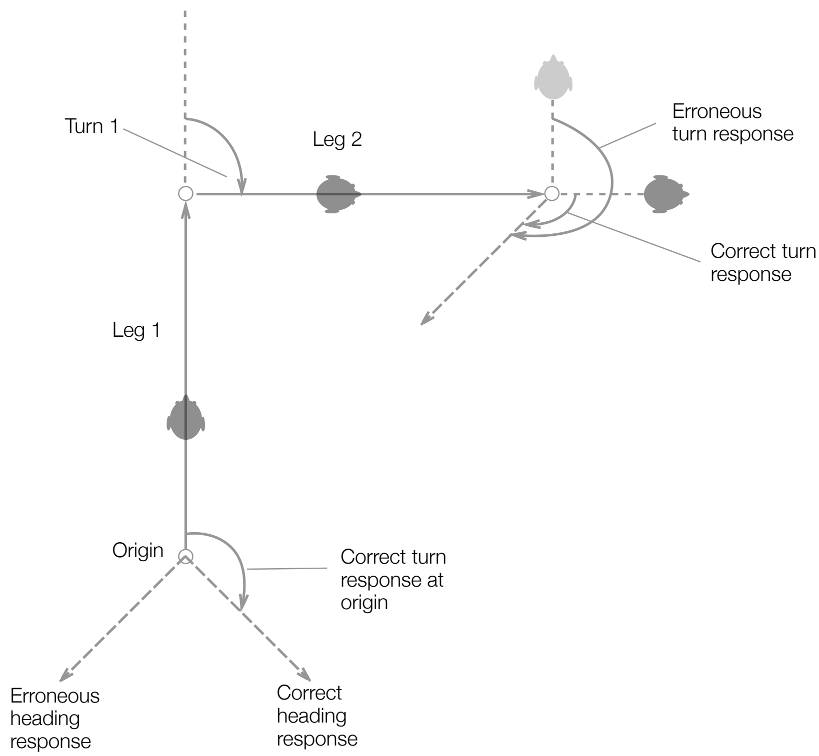

Figure 1. Schematic of the triangle-completion task used by Klatzky et al.

Note. The participant was presented with the path consisting of Leg 1, Turn 1, Leg 2 and then instructed to turn and face the origin. Participants who did not update their heading (indicated by a faded head) erroneously overturned by the value of Turn 1. Figure adapted from the original schematic in Klatzky et al. (1998).

Five conditions were used to investigate the effect of different stimuli on the spatial updating process, described in Table 1.

Table 1. Experimental conditions in Klatzky et al. **(1998)**

Condition	Description
Walk	Participants were blindfolded and led along the first leg, then turned, then led over the second leg.
Describe	Participants were blindfolded and then heard a verbal description of the pathway, with leg lengths described in meters and turns described in degrees.
Watch	Participants viewed the experimenter walking the pathway, and then closed their eyes before responding.
Visual turn	Participants sat on a stool, viewing a virtual environment via head-mounted display. The environment depicted a field of vertical posts resting on the ground plane. In each trial, participants watched a visually simulated movement according to the pathway, before responding.
Real turn	Same as Visual turn, but during the turn, participants were physically rotated on a chair so that the physical rotation matched the visual one.

Note. The five experimental conditions in a seminal study by Klatzky et al. (1998), investigating how participants updated their mental representation of space during real, imagined and simulated movement.

The researchers found that when participants did not turn physically (in the Describe, Watch and Visual turn conditions), heading errors increased with the turning amount. Suggesting that participants were using some kind of cognitive strategy, not automatic spatial updating. However, when participants did turn physically (in the Walk and Real turn conditions) heading errors did not correlate with turning amount.

These findings have lent support to the notion that allowing participants to physically perform simulated movements enables spatial updating. In addition to this, studies have found that when physical motion cues are missing, spatial updating seems to be impaired (Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004).

In other words, the prevailing notion is that for spatial updating to trigger automatically, physical motion cues are necessary (Chance, Gaunet, Beall, & Loomis, 1998; Ruddle & Lessels, 2006; Wraga et al., 2004). When they are present, spatial updating is facilitated (Farrell & Robertson, 1998; Presson & Montello, 1994; Rieser, 1989). When they are not, spatial updating is impaired (Klatzky et al., 1998; May & Klatzky, 2000; Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004).

Challenging notion: No, not always

The notion that physical motion cues are required to automatically trigger spatial updating has been challenged in recent studies however. Using a rapid pointing task and highly structured photorealistic replica of familiar natural environments, Riecke and colleagues showed that visual cues alone can be sufficient for automatic and obligatory spatial updating (Riecke & Bülthoff, 2004; Riecke, Heyde, & Bülthoff, 2001). In these studies, concurrent physical motion cues showed little, if any, effect on the quality of spatial updating, measured as a function of response time and pointing accuracy.

What could explain this apparent conflict? Klatzky et al. (1998) and others used optic flow for visuals, where participants could detect movement but no discernible features – similar to navigating in a snowstorm – while Riecke et al. used highly structured and photorealistic visual stimuli. The experiment suggests that the properties of the visual stimulus itself play an important part when it comes to triggering automatic spatial updating.

In a follow-up study, Riecke et al. (2007) found that a natural and structured visual scene did indeed suffice to enable automatic and obligatory spatial updating, irrespective of concurrent physical motions. Furthermore, displaying optic flow devoid of landmarks during motion and pointing phases proved to be insufficient to trigger automatic spatial updating, even when physical motion cues were added.

To summarize the relevant research, and attempt to answer the question posed in the above heading: Physical motion cues have been found to facilitate automatic spatial updating during rotations. Additionally, the absence of these cues can impair automatic spatial updating, but it seems that this is not always the case, depending on the structure of visual stimuli and other potential factors.

Regardless of whether physical motion cues are necessary to trigger automatic spatial updating or not, we know that they play an important part in this process. However, we are still unsure how important they are compared to the visual stimulus. In order to address that question, we will delve further into the theory of how our brains represent spatial information.

2.3.2. Egocentric, allocentric or both?

There are two main models of how spatial information is represented in the brain: The *egocentric model*, where the reference frame is centered on the observer, and the *allocentric model,* where the reference frame is not centered on any specific observer or location but all of them at once. To simplify, the egocentric model can be imagined as a first-person view of an environment, whereas the allocentric model can be likened to a top-down map view (Klatzky, 1998).

Initially, research focused on confirming the exclusive existence of one representation over the other, but more recent findings suggest that both egocentric and allocentric representations exist in parallel, and that they combine to support behaviour depending on the task (Burgess, 2006). How we choose which representation to use at which time is still unclear, but there is evidence that the choice is dependent on the sources of information provided (visual, vestibular, proprioceptive) (Thinus-Blanc & Gaunet, 1997) as well as environmental factors like structure, familiarity of the scene, and the amount of movement (Burgess, 2006).

To complicate things further, it seems that personal preference also plays a part in choosing a reference frame to use. In a study where participants indicated their end position after a virtual navigation task, Gramann et al. (2005) found that some participants reacted as if they had taken on the new orientation during turning, whereas others consistently failed to update their heading while updating their position accurately. Many other studies have confirmed these findings (Goeke, König, & Gramann, 2013; Gramann, 2013; Gramann et al., 2010; Gramann, El Sharkawy, & Deubel, 2009; Gramann et al., 2006; Gramann, Wing, Jung, Viirre, & Riecke, 2012; Riecke, 2008, 2012).

Gramann et al. argued that the participants who updated their heading correctly did so because they preferred to use an allocentric reference frame to update their spatial representation, and categorized them as "Non-Turners" (see Figure 2 for an illustration). Participants who did incorporate heading changes however, were believed to prefer an egocentric reference frame, and were categorized as "Turners". Furthermore, they found that pure visual information without physical motion cues can suffice to build up an egocentric spatial representation, i.e. Turner behaviour, thus challenging the view that physical motion cues are required for developing an egocentric representation (Gramann et al., 2005).

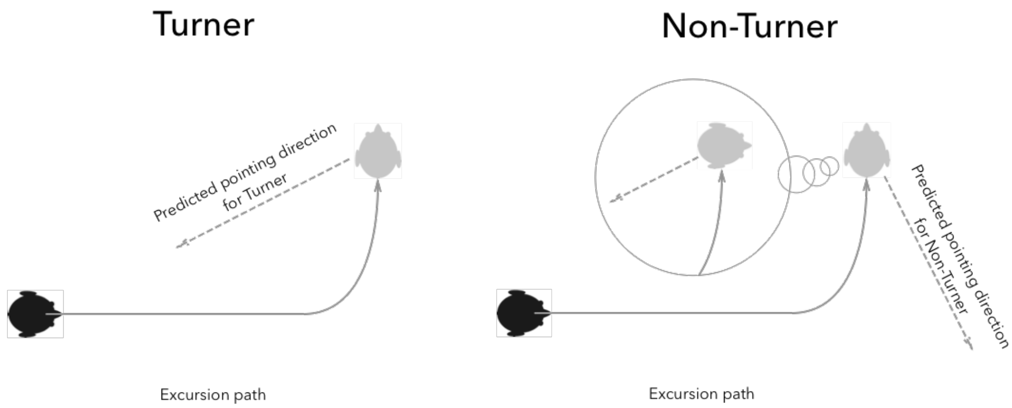

Figure 2. Explanation of Turner vs. Non-Turner behaviour

Note. **Left:** A Turner pointing response after travelling along a path with a 90-degree turn to the left. Having correctly updated their heading during the excursion, the participant points accurately towards the starting point. **Right:** A Non-Turner response after travelling the same path. As heading fails to update, the participant points to where the starting point would be if they had not rotated at all during the excursion (depicted inside thought bubble).

In addition to finding that participants had a distinct preference for one reference frame over the other, Gramann et al. found that when instructed to use the non-preferred one, both groups displayed no decline in response accuracy relative to their preferred reference frame (Gramann et al., 2005). These findings lend further support to the assumption that egocentric and allocentric spatial representations coexist in the brain during navigation, and that they work in parallel even though one may be preferred over the other.

Why do some people prefer to use different spatial strategies, and what are the differences in the resulting representations? This question remains largely open, although recent studies have shown that participant factors such as gender and ethnicity affect this preference. Goeke et al. (2013) found female participants to predominantly use a Non-Turner strategy when asked to select one out of four homing arrows to indicate the initial starting location, after watching a video of virtual passages through a star-field with one turn in either the horizontal or vertical axis. Male participants however, used both Turner and Non-Turner strategies with comparable probabilities. In a similar task, Sproll (2013) found ethnicity to influence preference for spatial strategies, with Caucasians showing a higher probability for Turner behaviour. In this study however, we will focus on the external factors that might play a key role in the choice of spatial strategy; namely visual and vestibular/proprioceptive motion cues.

3. Experiment 1: Virtual City

3.1. Motivation and research goal

As detailed in the previous section, two conflicting notions exist to explain the role of physical motion cues in enabling automatic spatial updating: The prevailing opinion that they are indispensable for proper spatial updating and spatial orientation (Bakker, Werkhoven, & Passenier, 1999; Chance et al., 1998; Klatzky et al., 1998), and the challenging notion that visual cues alone can be sufficient for automatic spatial updating (Riecke et al., 2007; Riecke, Heyde, et al., 2005).

In this experiment, our main motivation was to build upon and extend this line of research. Using detailed photo-realistic replica of real scenes rather than optic flow (a moving visual field without salient landmarks, like a star field), Riecke et al. found that visual information alone could indeed suffice to elicit automatic and obligatory spatial updating during rotations in a virtual environment (Riecke et al., 2007; Riecke, Heyde, et al., 2005). However, they did not test translations in their studies. This experiment complements these findings by adding a translational factor to further investigate the contribution of physical motion cues to automatic spatial updating. When moving along a path in a virtual city with naturalistic and highly structured visuals, are physical rotations still needed to enable automatic spatial updating, or can the visuals alone suffice?

To this end, we used a virtual spatial updating task based on an established point-to-origin paradigm (Gramann et al., 2005; Klatzky et al., 1998; Riecke, 2008). Participants moved passively along streets of varying curvature, at the end of which they were asked to point back to the origin of the path using a modified joystick. To implement physical rotations, we seated participants on a software-controlled rotating chair.

We purposefully did not include any salient landmarks. This was so that participants could not solely use piloting (landmark-based navigation) to perform the task, but had to incorporate rotational and translational cues (path integration). Therefore, our study did not require participants to establish a “cognitive map”, although some might have (Ishikawa & Montello, 2006).

3.2. Hypotheses

We hypothesized as follows:

H1: Passive physical motion cues enable automatic spatial updating

Based on previous blind walking and optic flow studies (Farrell & Robertson, 1998; Rieser, 1989; Ruddle & Lessels, 2006), adding passive physical motion cues should enable automatic spatial updating. This in turn should yield improved point-to-origin performance (unless the visual cues were already sufficient to enable automatic spatial updating). This would indicate that physical motions are necessary for successful spatial orientation in VR, even if naturalistic visuals are used. Conversely, should performance not improve when adding physical motions, we would conclude that either the visual environment was sufficient to fully trigger automatic spatial updating, or the physical cues failed to trigger spatial updating.

H2: Pointing errors increase for larger turning angles

Based on research on imagined perspectives switches we expect larger pointing errors for increasing turning angles (Klatzky et al., 1998; Riecke, 2008; Wiener & Mallot, 2006).

H3: Pointing takes no longer for larger turning angles

If spatial updating is automatic, no additional processing time should be needed at the end of the motion, as the mental representation was already automatically updated during the motion (Farrell & Robertson, 1998; Rieser, 1989). That is, response times should show little if any increase with turning angle. Conversely, should the response time clearly increase with turning angle, we could infer that spatial updating was not fully automatic.

H4: We do not observe Non-Turner behaviour in the VISUAL ONLY TURN condition

If the conditions where only visual motion cues are provided (no physical motion) are sufficient to trigger obligatory spatial updating, rotations should always be updated during those conditions. That is, we should only observe Turner behaviour. Hence, any observation of consistent Non-Turner behaviour would indicate that the visual cues alone were insufficient to obligatorily trigger spatial updating, at least for those participants.

H5: We do not observe Non-Turner behaviour in the REAL TURN condition

Similarly, if added physical rotations trigger obligatory spatial updating as predicted by the literature (Klatzky et al., 1998; Presson & Montello, 1994; Rieser, 1989), we should not observe any failures to update rotations during those conditions. That is, we should only observe Turner behaviour.

H6: Passive physical motion cues obligatorily trigger spatial updating

Conversely, observing less Non-Turner behaviour in the REAL TURN condition would indicate that adding passive physical rotation cues facilitates spatial updating.

3.3. Methods

3.3.1. Participants

A total of 12 Simon Fraser University undergraduate and graduate students (4 female) voluntarily participated in all parts of this study. They either received monetary compensation at standard rates, or course credit. Participant ages ranged from 18 to 33 years (mean = 22.5 years).

All participants had normal or corrected-to-normal vision, reported no history of motion sickness and were naïve to the purposes of the experiment. The study was approved by the university's ethics board. Written informed consent was obtained from each participant before the experiment.

3.3.2. Stimuli and apparatus

The virtual environment was displayed non-stereoscopically using an eMagin Z800 3D Visor HMD at a resolution of 800 x 600 pixels and a field of view of 32° x 24° at 60 Hz. Head movements were tracked via a Polhemus 6 degree-of-freedom motion tracker. Participants wore active noise-cancelling headphones and a blindfold mask over the HMD to exclude all auditory and visual cues from the surrounding lab.

Participants were seated on a chair mounted centrally on a 2 x 2m computer-controlled motion simulator, as illustrated in Figure 1 (see http://iSpaceLab.com/iSpaceMecha). The virtual environment was created using Procedural's CityEngine 3D modeler and rendered using Worldviz Vizard software. It consisted of a three-dimensional model of a city environment that contained ten individual curved street segments, surrounded with buildings (see Figure 1). Each street segment was designed to have a 45m long straight portion followed by a 40m long curve of 10°, 50°, 90°, 130° and 170° in either direction (see Figure 3 for an illustration of the setup). Although naturalistic, the virtual scene did not contain any salient landmarks that participants could have used for determining where they were relative to the starting point.

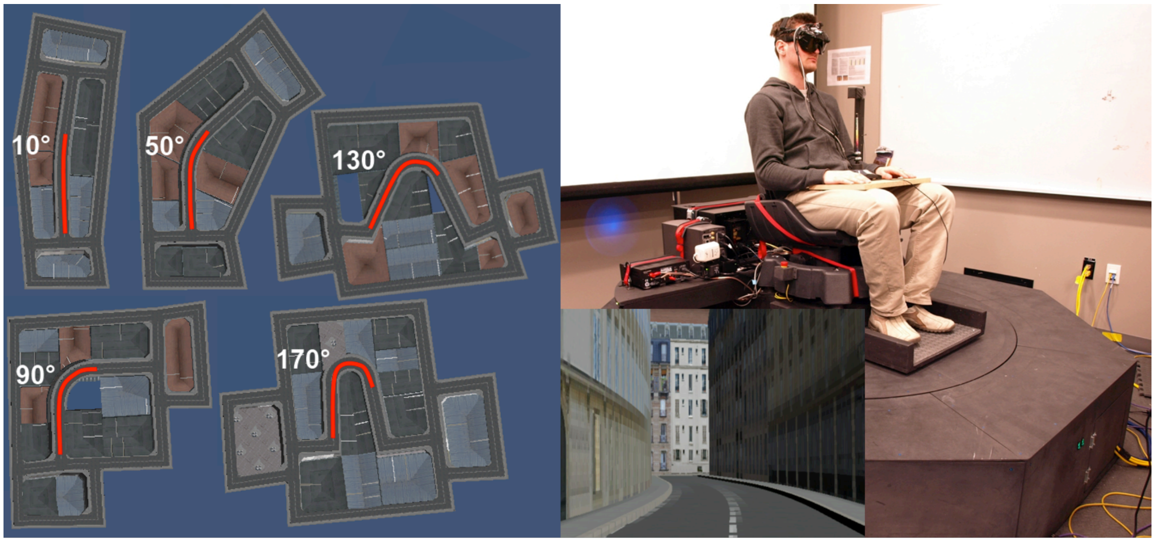

Figure 3. The experimental setup

Note. Left: A top-down view of the curved street segments used in the experiment (left-turn counterparts not shown). Right: Participants sat in a chair on a circular motion simulator, viewing the virtual environment through a head-mounted display. Bottom inset: The starting point of a 10-degree left-turn trial from the participant's viewpoint.

3.3.3. Procedure

Each trial involved a passive motion phase and a pointing phase. The motion phase consisted of a translation and rotation along one curved street segment within the virtual environment (3 m/s maximum translational velocity with a short acceleration and deceleration phase, 40°/s maximum rotational velocity with an acceleration of $50°/s^2$).

Upon arriving at the end of the trajectory, participants were asked to point "as quickly and accurately as possible" to the origin of the movement as if they had physically traveled it. Participants pointed with a modified Logitech Attack 3 joystick that was mounted on a wooden board and positioned on the participant's lap.

Two rotation conditions were compared: In the REAL TURN condition, participants rotated on the motion simulator as their viewpoint rotated in the virtual environment. In the VISUAL ONLY TURN condition, participants did not physically rotate. A real-world practice phase was used to ensure that they understood the procedure and could consistently point with at least 20° accuracy to a visible target. There was no visual indication of pointing response, so participants had to rely on proprioceptive and haptic cues to indicate in which direction they were pointing.

Participants never received any feedback on pointing accuracy. This was done to prevent participants from using cognitive strategies or recalibration for the pointing task, as previous studies have shown that when given unlimited response time and feedback participants can perform point-to-origin tasks relatively well (Wiener & Mallot, 2006).

3.3.4. Experimental design

We used a 2 (rotation condition: REAL TURN, VISUAL ONLY TURN) x 2 (turning direction: left, right) x 5 (turning angle: 10°, 50°, 90°, 130°, 170°) within-participant experimental design. The main experiment had 3 sessions, consisting of 20 trials each (10 for each of the 2 rotation conditions in balanced order). Rotation conditions were blocked within each session, while the virtual turning direction and angle were randomized. The experiment took less than one hour overall. See Figure 4 for a detailed diagram explaining the experimental procedure.

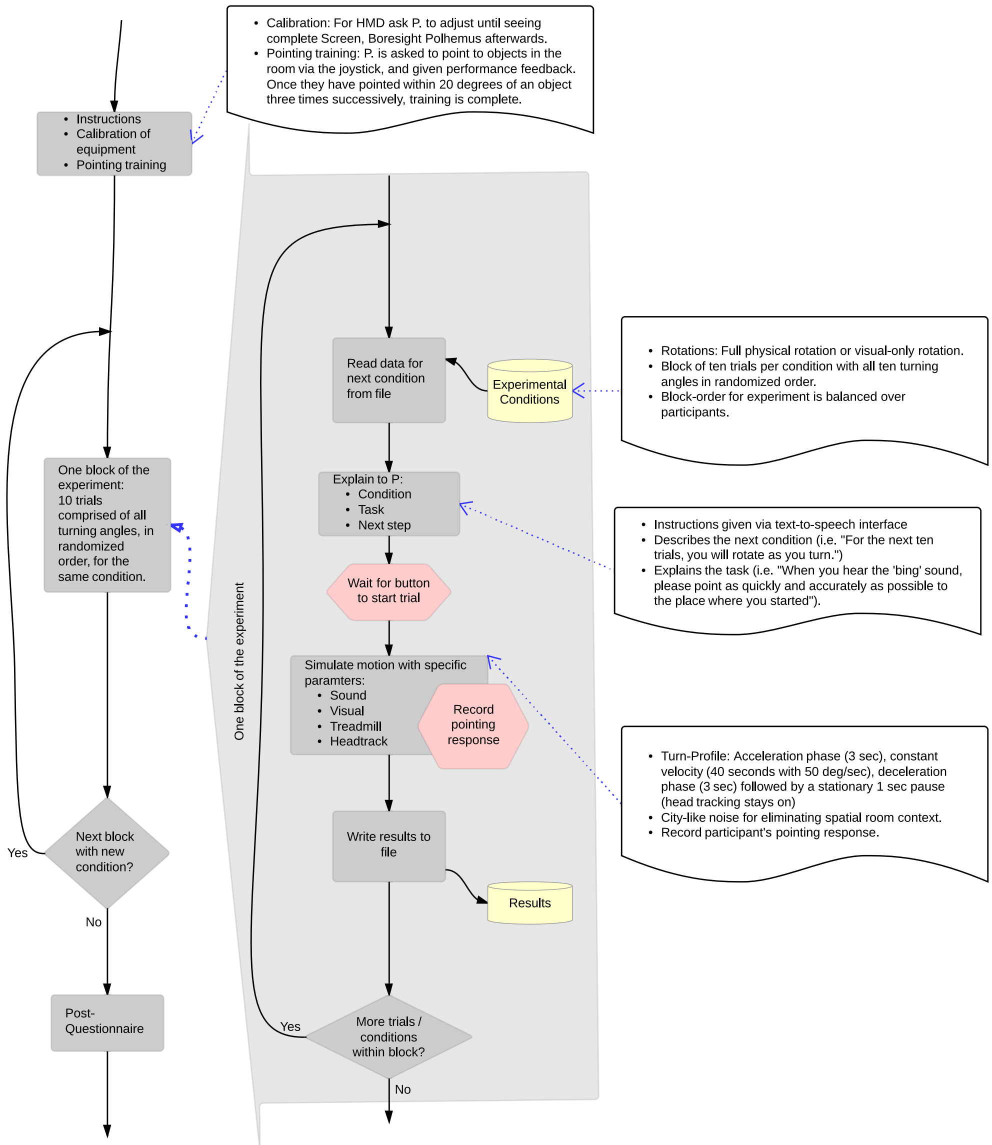

Figure 4. Diagram explaining the procedure in Experiment 1

3.4. Results

Pointing data were pooled over the left and right turning directions (which were not the focus of the study), and analyzed using two-way repeated measures ANOVAs.

Greenhouse-Geisser correction was applied where needed. See Table 2 and Table 3 for an overview of independent and dependent variables, respectively.

Table 2. Name, description and range of independent variables used in Experiment 1

Variable name	Description	Range
Rotation condition	Whether or not the participant rotated physically during the trial.	Binary: VISUAL ONLY TURN or REAL TURN
Turning angle	Turning angle of the path travelled during the trial.	Five levels: 10, 50, 90, 130 or 170 degrees.
Turning direction	Turning direction of the path travelled during the trial.	Binary: LEFT or RIGHT.

Table 3. Name, description and range of dependent variables used in Experiment 1

Variable name	Description	Range	Data type
Absolute pointing error	The absolute difference, in degrees, between the pointing direction and the correct homing direction. Averaged across trials.	0 – 160 degrees.	Ratio
Response time	Time passed, in seconds, from the moment a participant was instructed to point until she had settled on a pointing direction and a pointing response was registered.	0.2 - 10 seconds.	Ratio

When analyzing pointing data, we observed that two participants showed consistent Non-Turner behaviour, as they always pointed as if they had not incorporated any rotations at all (Participants 3 and 5 in Figure 5). To prevent these qualitatively different responses from distorting the analysis of remaining participants, we separated them from the main group during analysis.

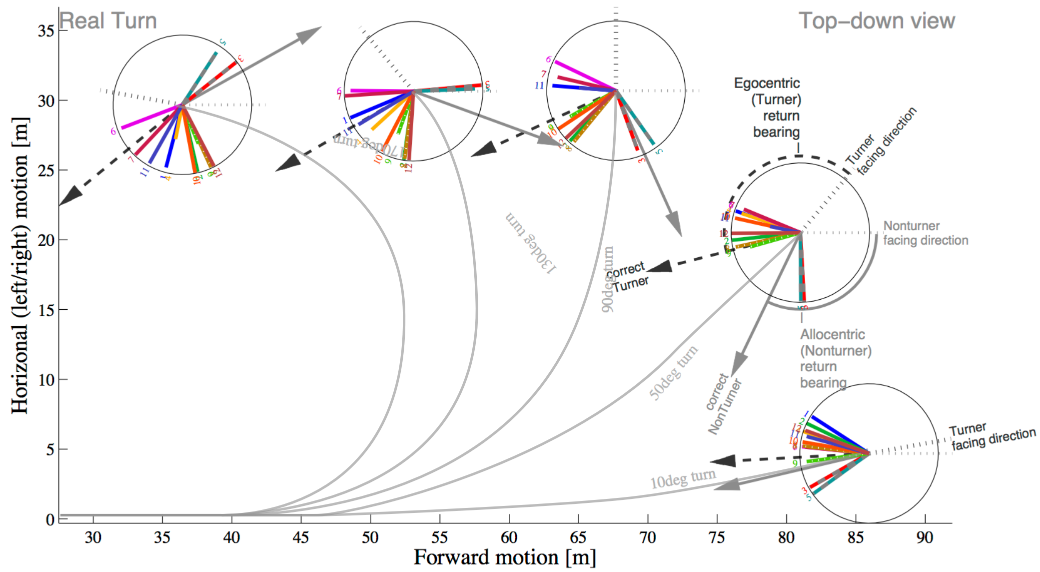

Figure 5. ***A visual representation of pointing responses in REAL TURN conditions.***

Note. From left to right, each circle represents a destination arrived at after traveling along a path with a 170, 130, 90, 50 or 10 degree turn respectively (averaging over left/right turns). Each coloured line represents the circular mean pointing response for one participant, with its radius measuring the concentration of angles (i.e. a short line indicates high variance in pointing responses). Correct pointing responses are shown with black dotted arrows, whereas gray solid arrows represent the pointing responses should heading changes fail to update during motion.

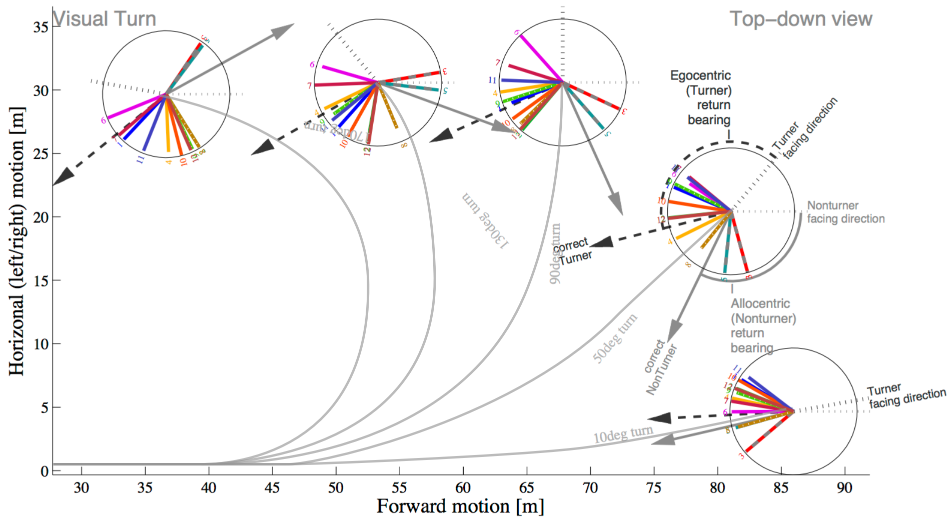

Figure 6. A visual representation of pointing responses in VISUAL ONLY TURN conditions.

Note. From left to right, each circle represents a destination arrived at after traveling along a path with a 170, 130, 90, 50 or 10 degree turn respectively (averaging over left/right turns). Each coloured line represents the circular mean pointing response for one participant, with its radius measuring the concentration of angles (i.e. a short line indicates high variance in pointing responses). Correct pointing responses are shown with black dotted arrows, whereas gray solid arrows represent the pointing responses should heading changes fail to update during motion.

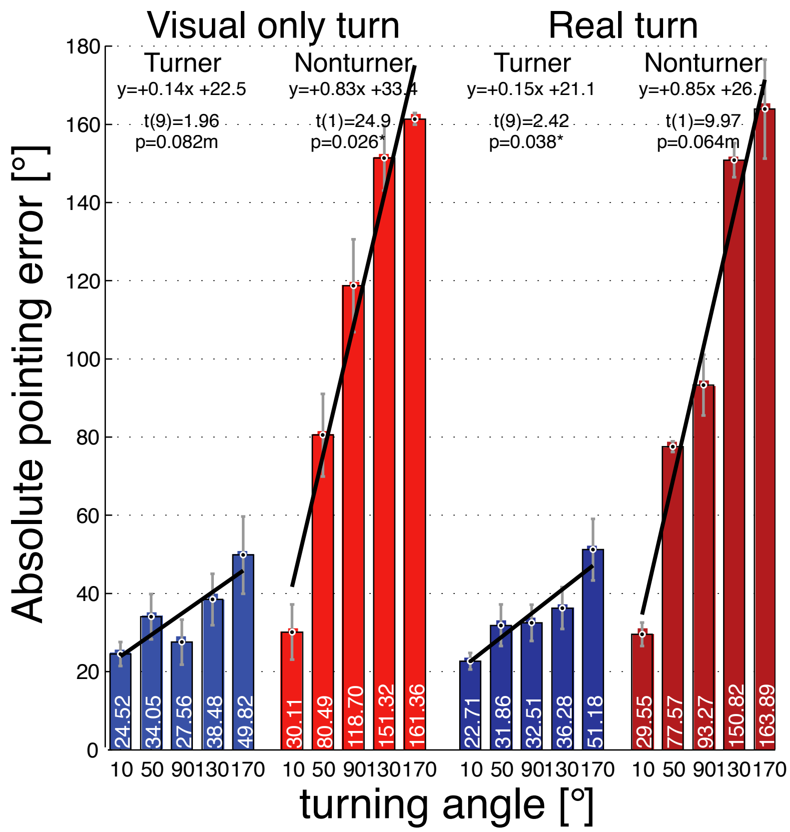

Figure 7. ***Absolute pointing error for each turning angle***

Note. Blue columns show the absolute pointing errors for participants exhibiting Turner behaviour, and red columns show the absolute pointing errors for the two participants exhibiting Non-Turner behaviour, in the Visual Only Turn and Real Turn conditions. Whiskers denote one standard error from the mean. Black lines represent a linear regression for each group, with slope equations and t-test results displayed above.

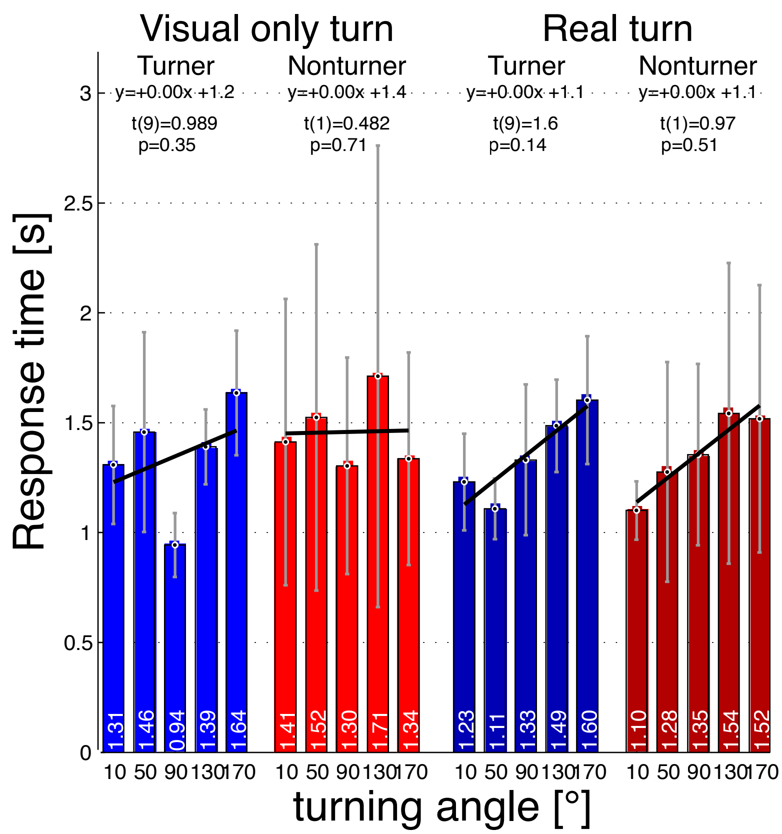

Figure 8. Response time for each turning angle

Note. Blue columns show response times for participants exhibiting Turner behaviour, and red columns show response times for the two participants exhibiting Non-Turner behaviour, in the Visual Only Turn and Real Turn conditions. Whiskers denote one standard error from the mean. Black lines represent a linear regression for each group, with slope equations and t-test results displayed above.

H1: Passive physical motion cues enable automatic spatial updating

The results, unexpectedly, showed no significant effect of turning condition on absolute pointing error, $F(1, 9) = .123$, $p = .734$, signed pointing error, $F(1,9) = .130$, $p = .727$ or response time, $F(1, 9) = 1.645$, $p = .232$. That is, we found participants to be no better spatially oriented when they received physical motion cues. According to our hypothesis we conclude that either the visual environment was sufficient to fully enable automatic spatial updating in all conditions, or that physical cues failed to enable automatic spatial updating. For whatever underlying reasons, adding physical motion cues to the current setup and procedure showed no benefits, which is noteworthy given past VR research.

H2: Pointing errors increase for larger turning angles

Turning angle significantly affected absolute pointing error, $F(1.526, 13.738) = 5.193$, $p = .028$. Although absolute pointing errors generally increased with increasing turning angles, Bonferroni-corrected post hoc tests showed no significant pairwise difference. This increase in pointing error is potentially due to accumulating path integration errors and/or higher task difficulty. However, as indicated in Figure 7, the slope of the linear regression fit is positive and different from 0 (marginally significant for visual-only turns and significant for real rotations), indicating that overall larger turning angles led to increasing absolute pointing errors (but not response times, see Figure 8).

H3: Pointing takes no longer for larger turning angles

Response time was on average 1.35s and was significantly affected by turning angle, $F(1, 9) = 2.693$, $p = .046$. However, correlations did not reach significance ($t(9) = 1.6$, $p = 0.14$ for REAL TURN, and $t(9) = 0.99$, $p = 0.35$ for VISUAL ONLY TURN), suggesting the ANOVA effects may have been spurious. Thus, on average, participants pointed neither faster nor slower as the turning angle increased. As this is one of the indicators of automatic spatial updating, this suggests that the visual cues may have been sufficient for enabling automatic spatial updating, irrespective of whether they were accompanied by matching physical rotations. Further studies with greater statistical power might help to address this, as detailed in Experiment 2.

H4: We do not observe Non-Turner behaviour in the VISUAL ONLY TURN condition

As mentioned above, careful analysis of all the experimental conditions revealed that 2 of the 12 participants (#3 and #5, see Figure 6) consistently exhibited Non-Turner behaviour throughout all trials in the VISUAL ONLY TURN condition. This observation of two consistent Non-Turners indicates that the visual cues alone were insufficient to trigger spatial updating, at least for those two participants.

H5: We do not observe Non-Turner behaviour in the REAL TURN condition

Based on (Farrell & Robertson, 1998) we expected that adding physical rotations should yield obligatory spatial updating. Thus, we should not have observed any Non-Turners in the REAL TURN condition. Surprisingly, however, we again observed consistent Non-Turner behaviour for the same two participants (see Figure 5), even when they physically rotated. Note that both Non-Turners exhibited this behaviour from the initial trials, which were REAL TURN condition in their case. Hence, we can exclude the possibility that they transferred their Non-Turner strategy from the visual-only condition.

H6: Passive physical motion cues obligatorily trigger spatial updating

Contrary to what we expected, Non-Turner behaviour was not reduced when physical rotations were added. This suggests that spatial updating was by no means more obligatory in the REAL TURN conditions.

3.5. Discussion and intermediate conclusions

Our findings lend support to the notion that visual cues alone can be sufficient to trigger spatial updating, provided that they are naturalistic. Although we did not include a pure optic flow condition, comparing our results with the most similar prior study (Riecke, 2008) shows smaller absolute pointing errors (34.9° vs. 50.8°) and smaller circular standard deviations (10.2° vs. 31.4°) for that naturalistic city environment used in this experiment, corroborating the results of Riecke et al. (2007; 2005).

Furthermore, we observed two participants that exhibited Non-Turner behaviour, that is, they responded as if they were still facing the initial direction even though they were aware of the path trajectory. This is noteworthy for two reasons: First of all, previous research found greater numbers of Non-Turners in visual-only conditions (Riecke et al. 40% (2008), Gramann et al. 50% (2005), Klatzky et al. 100% (1998)). In this experiment the percentage of Non-Turners was considerably smaller (17%). We posit that this might, at least in part, be explained by the more naturalistic visual cues used, even though they contained no salient landmarks. This is promising for VR simulations in that is suggests that by further increasing display quality we might be able to fully prevent Non-Turner behaviour, which is an important step towards effective yet affordable VR.

Secondly, we found that the Non-Turner participants in this experiment continued to exhibit the same behaviour even when additional physical motion cues were provided. As described in our publications discussing these findings (Riecke, Sigurdarson, & Milne, 2012; Sigurdarson, Milne, Feuereissen, & Riecke, 2012), this is the first time that Non-Turner behaviour was reported despite physical motion cues and naturalistic visuals. This contradicts previous research which suggests that physical motion cues are sufficient to trigger obligatory spatial updating (Farrell & Robertson, 1998).

However, spatial cognition research tends to create more questions than it answers, and this experiment was no exception. Why did two participants exhibit Non-Turner behaviour even when they received physical motion cues? Did the naturalistic visuals render physical motion cues unnecessary, or did the physical motion cues fail to improve spatial orientation for other reasons? How might participants fare in a similar task in which the visual stimulus is sparse rather than structured and naturalistic, or some mixture of both, with sparse yet minimally structured visuals? These are some of the questions that motivated us to design a follow-up to this experiment, detailed in the next section.

4. Experiment 2: The Thin White Line

4.1. Motivation and research goal

Among noteworthy findings in Experiment 1, two stand out when contrasted with the literature: First, we found participants to be no better spatially oriented when they received physical motion cues, compared to no physical motion cues, indicating that physical motion cues alone may not suffice to enable automatic spatial updating in a virtual point-to-origin task. This lends support to the challenging notion posited by Riecke et al. (2007; 2005), by extending the path trajectory to include translations as well as rotations. However, we only tested one virtual environment – a naturalistic and structured city – making it difficult to discern whether the visual stimulus was "good enough" to render physical motion cues unnecessary, or whether the physical motion cues failed to enable automatic spatial updating for other reasons. That is, would physical motion cues continue to show no significant benefit even if the visual quality of the virtual environment was reduced to its bare minimum? Second, we observed two participants that consistently exhibited Non-Turner behaviour throughout all trials, even when physical motion cues were available, in contradiction to previous research suggesting that physical motion cues are sufficient to enable obligatory spatial updating (Farrell & Robertson, 1998).

In Experiment 2, our goal was to extend and build upon these findings, as well as test if we could replicate these findings with a different set of naïve participants. We had participants perform a point-to-origin task similar to the one used in our previous experiment, with the following adjustments:

Optic flow environment: Results from our previous experiment suggested that there might exist a "winner takes all" effect, where physical motion cues do not further improve performance if the visuals are rich and naturalistic enough. To investigate this further, we added a condition with low-fidelity optic flow visuals, where visual information alone should not suffice to enable spatial updating for movement trajectories including rotations (Chance et al., 1998; Klatzky et al., 1998; Riecke et al., 2007; Wraga et al., 2004).

Minimally structured optic flow environment: Riecke et al. (2004; 2001; 2005) showed that visual cues alone can suffice to enable automatic spatial updating, using highly structured photorealistic replica of familiar natural environments (see subsection 2.2.1 for a detailed discussion). To isolate and investigate the effects of structure, we added a condition with low-fidelity optic flow visuals that had been augmented to include a minimal amount of structure: a thin white line predicting the participant's trajectory without serving as a landmark (see Figure 12).

Categorization of pointing responses: In Experiment 1, two out of twelve participants pointed as if they had updated their position but ignored changes in heading, a type of spatial behaviour that has been described as Non-Turner (Gramann et al., 2005). We observed this behaviour even under conditions were physical motion cues were available – for the first time to the best of our knowledge.

To expand on these findings, we designed this experiment to investigate the spatial updating process specifically through the lens of Turner vs. Non-Turner behaviour. Rather than analyzing pointing responses in terms of their absolute error from the correct response and categorizing the participant as exhibiting Turner or Non-Turner behaviour – as we did in Experiment 1 – we devised a simple method to categorize a single pointing response as indicating either Turner or Non-Turner behaviour: If the participant pointed in the overall correct hemisphere (e.g., into the left hemisphere for a left turn), that pointing response was categorized as a Turner response. If they pointed in the overall incorrect hemisphere (e.g., into the right hemisphere for a left turn), that pointing response was categorized as a Non-Turner response (see Figure 9 for a visual explanation).

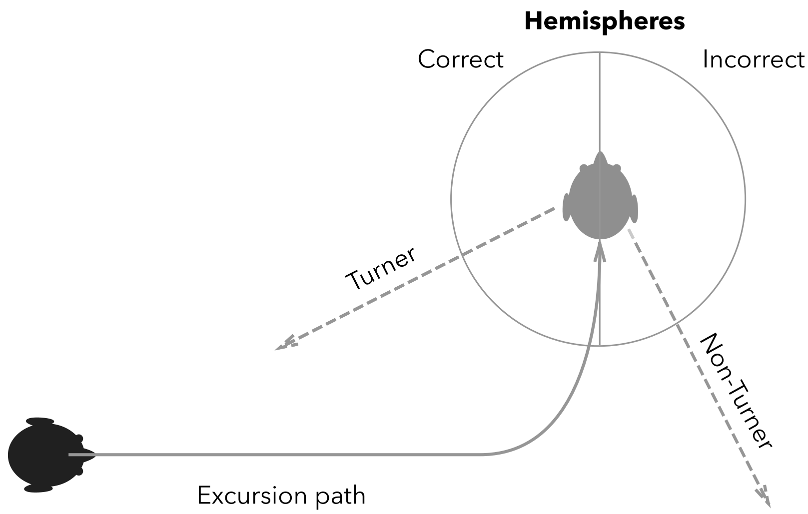

Figure 9. ***Description of pointing response categorization as Turner or Non-Turner based on overall hemisphere of pointing***

We opted for this novel (to the best of our knowledge) approach in an attempt to mitigate the inherent ambiguity of using absolute pointing errors to discern between qualitatively different pointing responses, and to place focus on pointing behaviour within each trial rather than categorizing participants overall as Turners or Non-Turners. In return, we run the risk of oversimplifying a complex process, however we believe that the benefits of this method can outweigh its limitations (see discussion on limitations in section 5.3).

Improved head-mounted display: In Experiment 1, participants wore a head-mounted display (HMD) with a resolution of 800 x 600 pixels and a field of view (FOV) of 32° x 24° at 60 Hz. As human spatial orientation in VR typically benefits from a large FOV (Péruch & Gaunet, 1998; Riecke, Heyde, et al., 2005; Riecke, Veen, & Bülthoff, 2002), we upgraded our equipment to use an HMD with a resolution of 1280x1024 pixels, and a FOV of 102° x 64°; a considerable improvement.

Improved pointing device: In Experiment 1, participants gave pointing responses by deflecting a modified joystick in the chosen direction. Although we found this method to be sufficiently accurate, we surmised that the mental transfer of a pointing direction to a matching joystick deflection might have been a confounding factor for participants during the experiment. We therefore devised a pointing device that afforded more accurate pointing (Haber, Haber, Penningroth, Novak, & Radgowski, 1993): A motion-tracked stick held with both hands, and pointed in the chosen direction when instructed (see Figure 14).

Difficulty ratings: In addition to measuring pointing responses, participants were asked to rate the difficulty of individual trials (termed Task difficulty) during the experiment, and the difficulty of each condition (termed Condition difficulty) after the experiment. The term “difficulty” was not explicitly defined for participants. This supplemental data can afford a richer analysis of the occurrence of spatial updating, as described in hypotheses H2, H4 and H9 below.

Increased statistical power: In this experiment, we increased the number of participants from 12 to 36, significantly empowering our statistical analysis.

4.2. Hypotheses

We designed Experiment 2 to test the following hypotheses:

H1: Added physical rotations increase Turner behaviour

If adding physical rotations enables spatial updating as the literature predicts (Chance et al., 1998; Klatzky et al., 1998; Wraga et al., 2004), we should see an increase in Turner behaviour during conditions with physical motion cues. Failure to observe this would lend confirmatory support to our findings from Experiment 1, and further support the notion that physical motion cues might not suffice for obligatory spatial updating, even when translations and rotations are combined in a smooth trajectory.

H2: Added physical rotations make the task feel easier

When performing a spatial task such as in this experiment, disoriented participants should find them more difficult than participants who maintain spatial orientation during the trial. Thus, by asking participants to rate the difficulty of a single trial, as well as a single condition, we can assess the occurrence and automaticity of spatial updating, as well as indirectly inferring the participant's cognitive load. If adding physical rotations triggers automatic spatial updating we should find that participants rate conditions with physical motion cues as significantly easier than conditions without. Failure to observe this would suggest that physical motion cues might not suffice to afford automatic spatial updating.

H3. Naturalistic and structured visuals increase Turner behaviour

According to Riecke et al. (2007; 2005), a naturalistic and structured visual stimulus can suffice to trigger automatic spatial updating even without any physical motion cues. Observing an increase in Turner behaviour under conditions with higher-fidelity visuals versus lower-fidelity would lend further support to that notion. Note that Riecke et al. used an abundance of landmarks in their scene, in order to render abstract cognitive strategies (like using symmetries and counting targets) virtually impossible. Here, we achieve the same goal by including no landmarks whatsoever. Furthermore, this experiment extends the tasks performed in Riecke et al. to include translations and rotations in a curvilinear path, rather than rotations alone.

H4: Naturalistic and structured visuals decrease task difficulty

Following the logic in H2, if a naturalistic and structured visual stimulus can suffice to trigger automatic spatial updating even without physical motion cues, we should expect participants to rate those conditions as easier than conditions with lower-fidelity visuals. Failure to observe this would suggest that naturalistic and structured visuals did not suffice to enable automatic spatial updating in the present study.

H5: Naturalistic and structured visuals decrease response times

If spatial updating is automatic, no additional processing time should be needed at the end of the motion as the mental representation was already automatically updated during the motion (Farrell & Robertson, 1998; Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004). Therefore, observation of significantly lower response times for a condition with naturalistic and structured visual stimulus would suggest that spatial updating was not fully automatic in other conditions.

H6: Larger turning angles do not increase response times

As above, when automatic spatial updating is active, no additional processing time should be needed at the end of the motion (Farrell & Robertson, 1998; Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004), so response times should show no significant increase with turning angle. Conversely, observation of response time increasing with turning angle would suggest that spatial updating was not fully automatic.

H7: Participants overestimate the magnitude of turning angles

In this experiment, we categorized pointing responses as Turner or Non-Turner, based on whether responses were in the overall correct hemisphere or not (i.e. rightwards for a right turn and vice versa). A Non-Turner response suggests that the participant failed to update their heading during the trial, but it could also result from a misperception of the turning amount (Riecke, 2008, 2012). That is, a participant who updates their heading correctly but significantly overestimates how far they have turned might be categorized as exhibiting Non-Turner behaviour in that trial, even if they correctly updated their heading during the trial. For example, a 170-degree right turn might be perceived as a 270-degree right turn, resulting in a pointing response towards the left hemisphere even if the participant correctly updated their heading.

How might we attempt to disambiguate these factors? A review of the literature shows that few studies have done so, allowing us to contribute a novel and potentially useful method. To separate these different factors of Non-Turner behaviour, we asked participants to draw an estimate of all turning angles they traveled during the experiment, in a post-experimental questionnaire (see Figure 10 for an example). We then measured the angle of the largest turn drawing, and compared to the largest actual turning angle used in the experiment (170 degrees). Participants who drew the largest turning angle as more than 180 degrees (allowing for a 10 degree error) were categorized as "over-estimators".

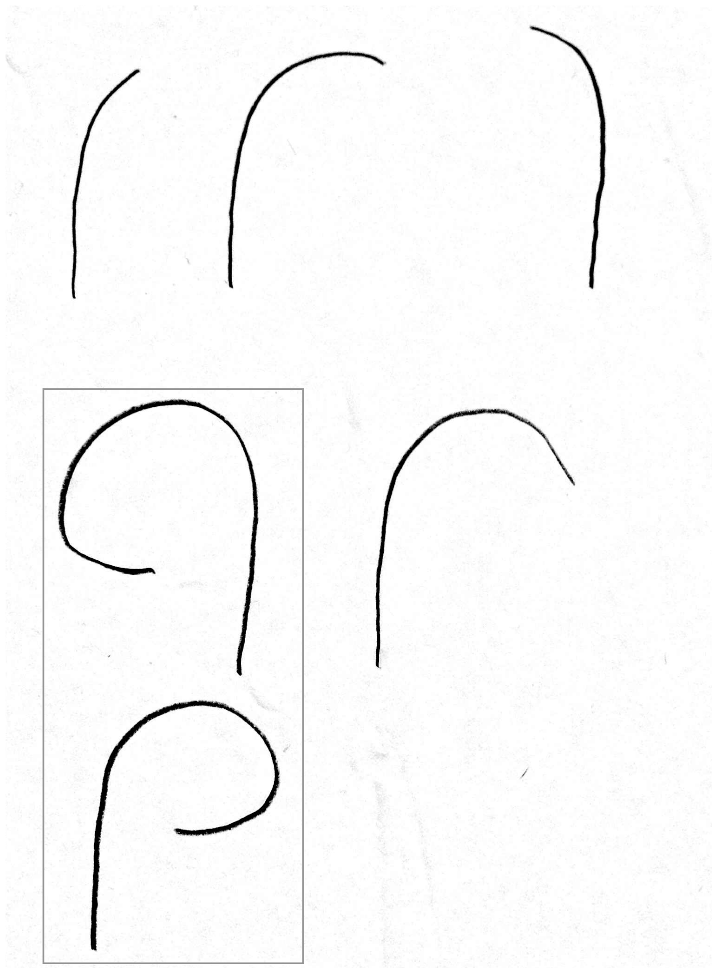

Figure 10. Illustration of participant drawings of path trajectories

Note. Out of all paths drawn by this participant, two represent the largest perceived angle (shown inside a box). In this case, the largest perceived angle by this participant is measured at 270 degrees, resulting in a categorization of “over-estimator”.

If this method of categorization is valid, we should see a marked difference in pointing behaviour between those who overestimated the largest turn, and those who did not.

H8: Participants “keep” the pointing strategies they adopted in the visual environment they experienced first

Many uses of Virtual Reality are based on the assumption that spatial knowledge acquired in a virtual world will transfer to the real world (see section 2.1.2 for a discussion on VR uses). Although real-world transfer was out of scope for this study, we wanted to investigate whether participants did transfer their spatial strategies between different visual environments. We therefore designed the order of each participant's exposure to visual environments to either go from lowest to highest fidelity or vice versa.

If participants do transfer their pointing strategies between visual environments, we should find that those who start in the high-fidelity environment exhibit lower Non-Turner rates across all visual conditions than those who start in the low-fidelity environment (assuming that higher-fidelity visuals elicit lower Non-Turner rates overall, see H3).

H9: Perceived difficulty of rotation conditions depends on the initially experienced visual environment

Following H8, we can also detect transfer effects of different visual environments on difficulty ratings for rotation conditions. This is of particular interest to us, as one of the underlying motivations for this thesis is the notion that physical motion cues do not further improve performance if the visual stimulus is rich and structured enough.

Our line of reasoning goes like this: If there exists a transfer effect of different visual environments, we can then measure the contribution of physical motion cues in the participant's mind by looking at how difficult they found conditions with physical motion cues in successive visual environments. That is, a participant who starts in a high-fidelity environment and partly transfers their strategy to the lower-fidelity environments may regard the added information from physical motion cues as redundant or even distracting, suggesting that their initial strategy did not benefit from physical motion cues. The opposite also holds true: A participant who transfers their strategy from low-fidelity to high-fidelity might find conditions with physical motion cues less difficult than those without.

H10: Mental rotation ability correlates with lower Non-Turner rates

Studies have shown that the ability to rotate mental representations of two-dimensional and three-dimensional objects correlates with increased performance in spatial tasks (Gardony, Taylor, & Brunyé, 2014; Malinowski, 2001). Therefore, observing that mental rotation ability is negatively correlated with the rate of Non-Turner trials would corroborate prior studies, and provide validation for our method of measuring spatial orientation.

4.3. Methods

4.3.1. Participants

A total of 36 Simon Fraser University undergraduate and graduate students (12 female) voluntarily participated in all parts of the study. They received $10 in compensation. Participant ages ranged from 19 to 37 years (mean = 25.1). All participants had normal or corrected-to-normal vision, reported no history of motion sickness and were naive to the purposes of the experiment. The study was approved by Simon Fraser University's ethics board, and written informed consent forms were obtained from each participant before the experiment.

4.3.2. Stimuli and apparatus

The virtual environment was displayed stereoscopically using an nVis SX111 Head-Mounted Display at a resolution of 1280x1024 pixels, and a field of view of 102x64 degrees at 60 Hz. Head movements were tracked via a Polhemus Liberty 6 degree-of-freedom motion tracker. To minimize external cues, participants wore active noise-cancelling headphones and the experimental room was dimmed during the course of the study.

Participants were seated on a chair mounted centrally on a 2x2m computer-controlled motion simulator, as illustrated in Figure 14. The virtual environment was created using Procedural's CityEngine 3D modeller, and rendered using WorldViz Vizard software.

The virtual environment consisted of three separate areas, representing progressively higher visual fidelity yet devoid of salient landmarks:

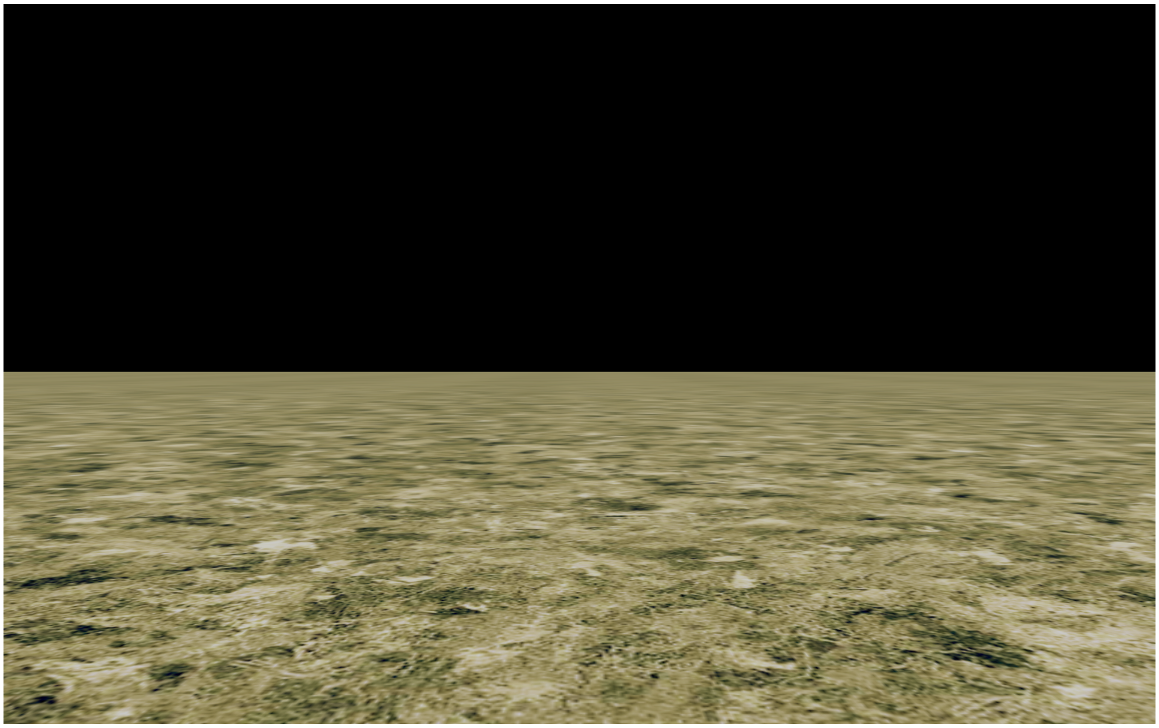

Figure 11. The Grass condition from the participant's viewpoint at the starting point

GRASS: A large ground plane with repeating textures, representing an optic flow environment. In this area, participants had neither landmarks nor a naturalistic scene to estimate their location relative to the starting point of movement (see Figure 11).

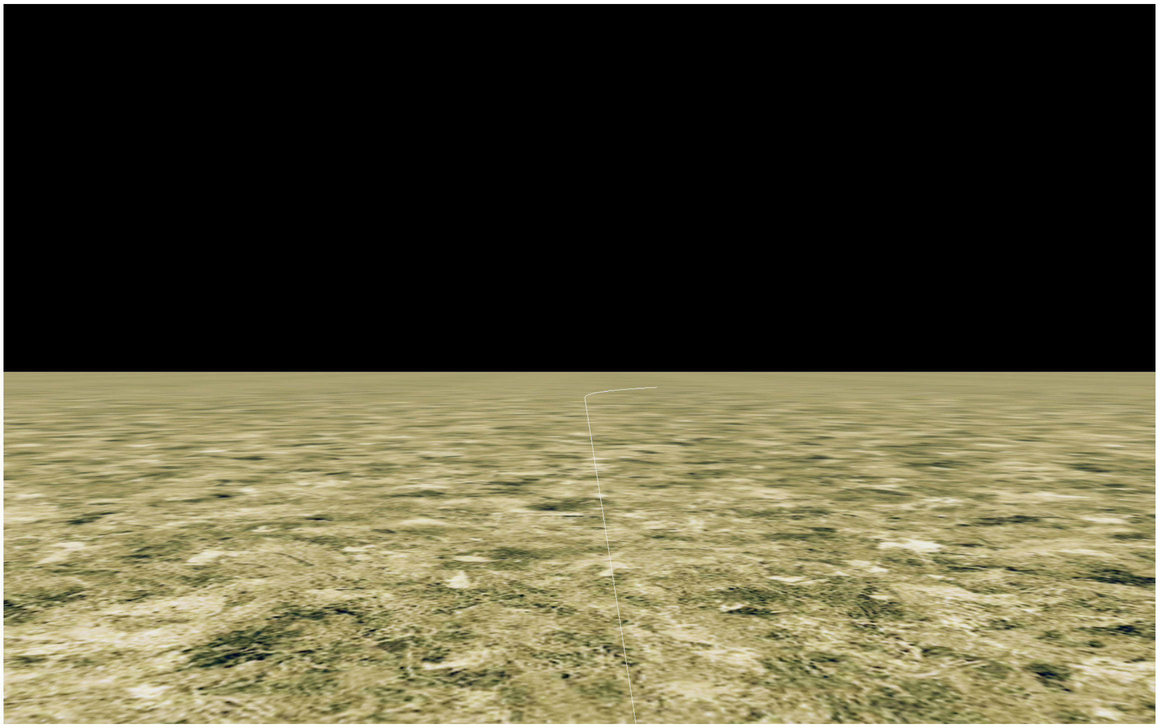

Figure 12. The GRASS WITH LINE condition from the participant's viewpoint at the starting point

GRASS WITH LINE: A large ground plane identical to the GRASS condition, with one difference: A thin white line was drawn on the ground plane, indicating the path trajectory for the current trial (see Figure 12). To prevent participants from seeing the entire line, and therefore enabling them to estimate the curvature of the path and thus calculate their position using cognitive strategies, an invisible blocking plane was applied to the line such that only a part of it could be seen at any time. In essence, the blocking plane served as an invisible row of houses on either side of the "street", maintaining a similar level of visual blocking as the participant experienced in the CITY condition.

This environment was designed to give participants a similar amount of optic flow as the GRASS environment, but with advance information about the direction of the upcoming turn without showing the total turning angle in advance. That is, participants still had to use path integration to estimate the overall turning angle, but had advance knowledge about the turning direction at the beginning of the trial.

Figure 13. ***The CITY environment from the participant's viewpoint at the starting point***

Note. This trial involves a 10-degree curve to the right.

CITY: A three-dimensional model of a city environment with ten individual curved street segments. Each street segment was designed to have a 45m long straight portion followed by a 40 m long curve of 10, 50, 90, 130 and 170 degrees in either direction. Although naturalistic, the virtual scene did not contain any salient landmarks that participants could have used to determine where they were relative to the starting point of their movement (see Figure 13). Similar to the GRASS WITH LINE condition, participants had advance knowledge about the turning direction at the beginning of the trial, but had to use path integration and updating of local features to estimate the overall turning angle, which was not discernible from the initial view.

Conditions

In this study, two rotation conditions and three visual conditions were compared. The two rotation conditions were REAL TURN and VISUAL ONLY TURN. In the REAL TURN condition, participants rotated on the motion simulator as their viewpoint rotated in the virtual environment. The physical rotation was performed at the exact same acceleration and speed as the virtual rotation. In the VISUAL ONLY TURN condition, participants did not physically rotate as their viewpoint in the virtual environment rotated. The three visual conditions were: GRASS, GRASS WITH LINE, and CITY, described above.

Participants never received any feedback on pointing accuracy. This was done to prevent participants from using cognitive strategies or recalibration for the pointing task, as previous studies have shown that when given unlimited response time and feedback, participants can perform point-to-origin tasks relatively well (Wiener & Mallot, 2006).

Figure 14. The experimental setup

Note. A participant sits in a chair on a motion platform, viewing the virtual environment (shown in top left) via head-mounted display, holding the pointing device with both hands.

4.3.3. Procedure

Before the experiment

Upon arrival at the lab, participants received a verbal and written description of the study, including information about potential adverse effects of simulator sickness and discomfort due to the head-mounted display. They then reviewed and signed a written form of informed consent. All participants were informed that they could stop the experiment at any time, for any reason, without affecting their compensation.

Prior to starting the experiment, participants completed a pointing training phase, where they pointed to various real-world objects in the experiment room, with a pointing stick that had a laser pointer affixed to the top, and calibrated so that the laser beam pointed in the exact same direction as the stick. Participants were instructed to point to items with eyes closed, and then open them to compare the position of the laser pointer to the object they pointed towards. This continued until participants could point accurately to items in the room, without having to readjust the stick.

During the experiment

The experiment consisted of six consecutive sessions, each corresponding to a combination of the two rotation conditions (REAL TURN and VISUAL ONLY TURN) and three visual conditions (GRASS, GRASS WITH LINE, and CITY) (see Figure 14 for an overview of the experimental setup). Each session consisted of 13 trials; three practice trials (wherein no data was recorded) followed by ten trials. The ten non-practice trials consisted of all combinations of the 5 turning angles (10, 50, 90, 130 and 170 degrees) and 2 turning directions (left, right), in randomized order. The three practice trials were random path segments chosen from the same pool of ten segments (see Figure 16 for a visual explanation of the experimental procedure).

Before each session, a computerized voice explained which visual and rotation condition the participant could expect during the next 13 trials, and that the first three trials would be regarded as practice. Upon completing a trial that involved a 50 or 130-degree turn, the participant was asked to indicate how difficult they found that trial, using the pointing stick to move a slider between 0 (extremely easy) and 100 (extremely difficult). We restricted task difficulty ratings to these two angles, rather than recording difficulty ratings after each trial, to minimize potential saturation effects of repeated ratings.

Each trial involved a passive virtual motion phase and a pointing phase. In the motion phase, participants moved passively along one curved street segment in the city area, or a curved path of the same length and curvature in the optic flow area. Regardless of area, the movement was made with a 3 m/s maximum translational velocity with a short acceleration and deceleration phase, and 40 degrees/second maximum rotational velocity with an acceleration of 50 degrees/second^2. Upon arriving at the end of the trajectory, participants were asked to point “as quickly and accurately as possible” to the origin of the movement, as if they had physically traveled it.

We used a 2 (Rotation condition: REAL TURN, VISUAL ONLY TURN) x 3 (visual condition: GRASS, GRASS WITH LINE, CITY) x 2 (turning direction: LEFT, RIGHT) x 5 (turning angle: 10, 50, 90, 130, 170 degrees) within-participant experimental design.

The main experiment had 6 sessions, corresponding with combinations of the three visual conditions and two rotation conditions. To balance the session order, participants were split into four gender-balanced, but otherwise randomized, groups that determined whether visual fidelity increased or decreased over the course of the experiment, and whether the first session included physical motion cues or not (see Figure 15 for an illustration).

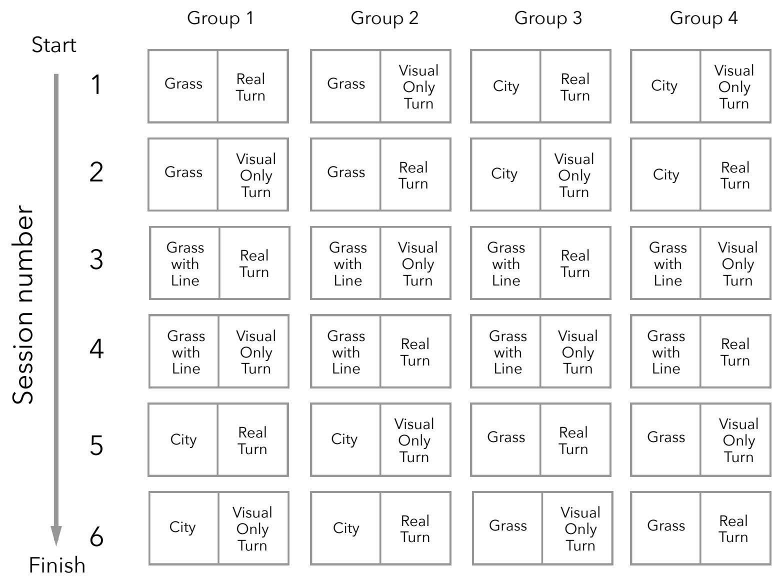

Figure 15. Balancing and order of conditions for Experiment 2

Note. Groups were of equal sizes (9 participants) and gender-balanced (6 males, 3 females). Note that groups 1 and 2 receive a progressively higher-fidelity visual stimulus over the course of the experiment, whereas groups 3 and 4 receive a progressively lower-fidelity visual stimulus.

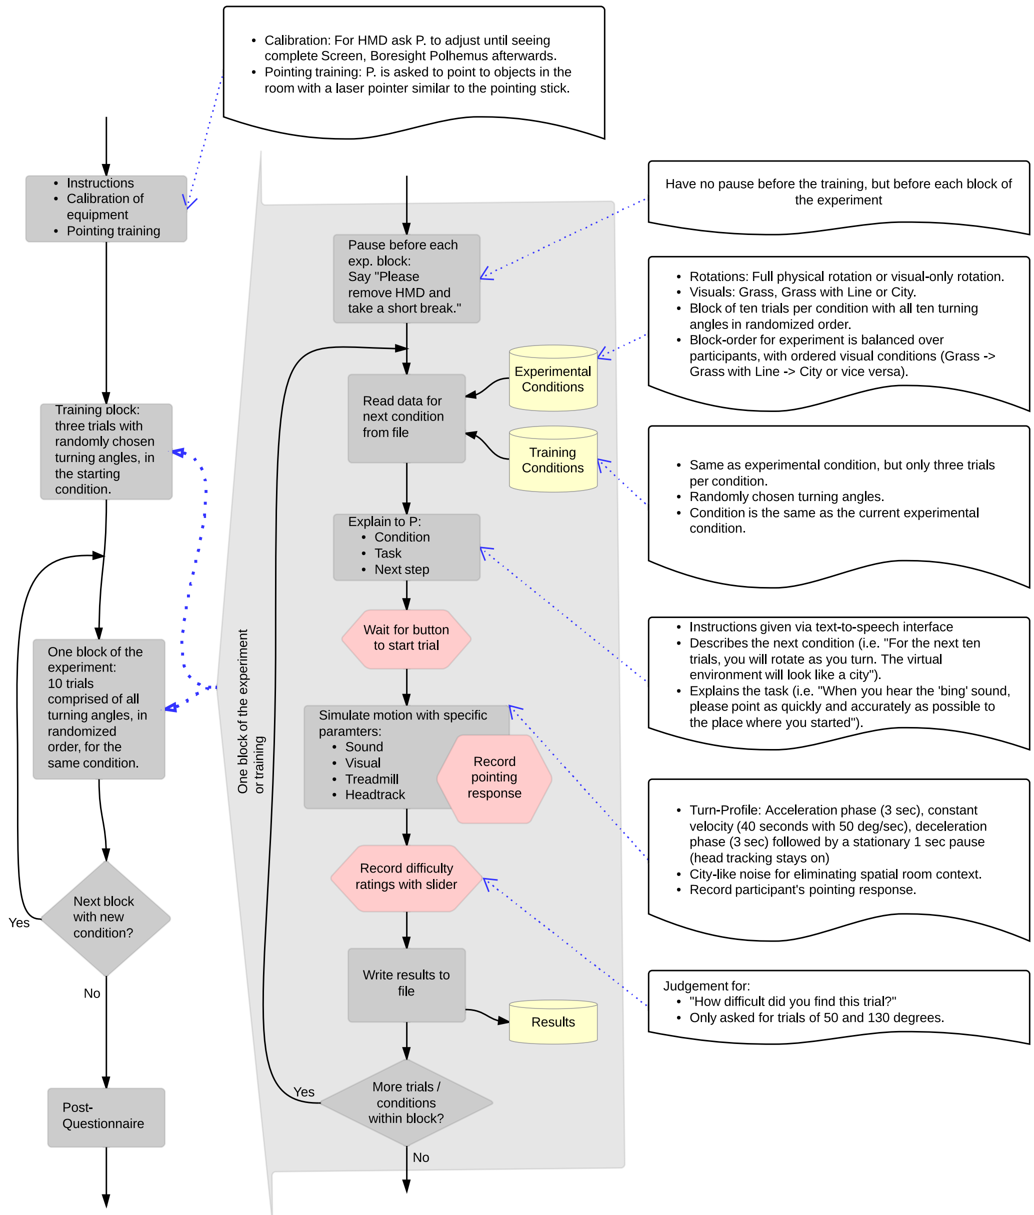

Figure 16. Diagram explaining the experimental procedure.

After the experiment

We used introspective measures to collect thoughts and feelings of participants after completing all trials, by way of a post-experimental questionnaire in four parts:

1. Condition difficulty ratings

Participants were given a sheet of paper containing an empty 2x3 table where cells represented the six combinations of stimuli they experienced during the experiment, and asked to give each one a difficulty rating between 0 (extremely easy) and 10 (extremely difficult).

2. Interview

Participants were interviewed briefly by the principal investigator, and asked to voice their thoughts and feelings on different parts of the experiment. The interview was semi-structured (see Table 4 for a list of questions asked).

Table 4. Post-experimental questions

1. How did you solve the task? Did you use any strategies?
2. What did you think of the physical motion used in the study?
3. Did you use the physical motion for spatial orientation?
4. What did you think of the visuals used in the study?
5. Did you use the visuals for spatial orientation?
6. Did anything bother you during the experiment?

During the interview, participants were also asked to rate their own every-day spatial orientation and sense of direction (on a scale from 0 to 10), and their visualization ability (on a scale from 0 to 10). For a detailed qualitative analysis of responses to these interview questions, see subsection 4.4.2.

3. Location memory test

In order to test for individual differences in spatial memory, and how these might relate to point-to-origin performance and the occurrence of Non-Turner behaviour, participants were asked to complete a location memory test after the interview (Silverman & Eals, 1992). Upon starting the test, participants were presented with a sheet of paper depicting 27 objects, which they were asked to inspect for one minute. The presentation array was then replaced with another sheet of paper, depicting the same 27 objects but with 14 of them in different positions. Participants then had one minute to indicate both objects that had moved, and those that had not moved (see Figure 17 and Figure 18).

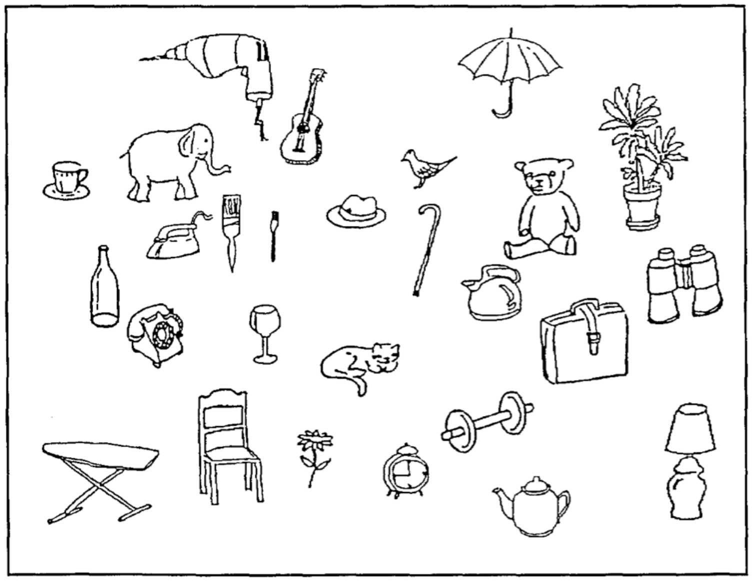

Figure 17. Presentation array of location memory test

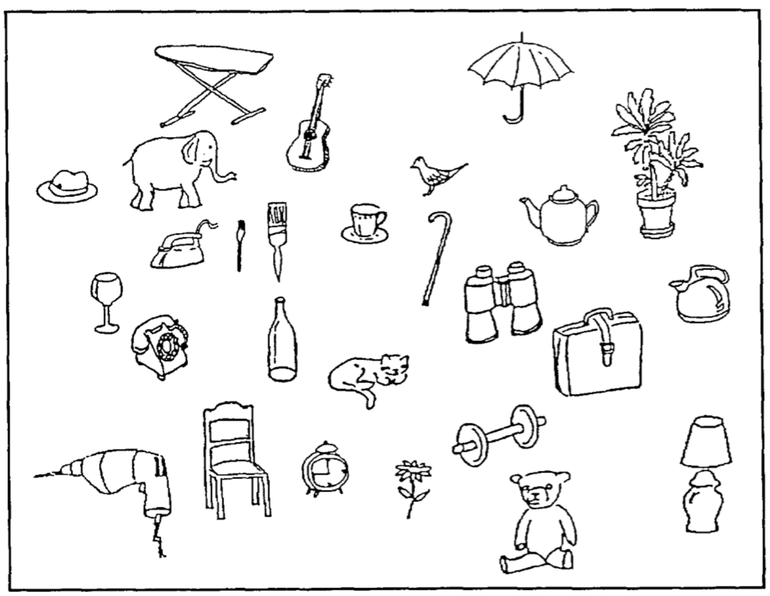

Figure 18. Response array of location memory test

4. Mental rotation test

In order to test for individual differences in mental rotation ability, and how these might relate to point-to-origin performance and the occurrence of Non-Turner behaviour, participants completed a refurbished version of the original Vandenberg & Kuse mental rotation test (1978), redrawn with help of a computer-assisted drawing program (Peters, Laeng, Latham, & Jackson, 1995). We chose this particular test, as it is one of the most commonly used measures of spatial ability.

The test is comprised of 24 items in which two-dimensional drawings of three-dimensional geometrical figures are to be compared. Each item consists of a row of five line drawings, including a geometrical target figure on the far left, followed by four response figures (see Figure 19). The participant's task was to indicate which two of the four response figures represented a rotated reproduction of the target figure. Participants were given 3 minutes to complete each subset of 12 items, separated by a 4-minute break.

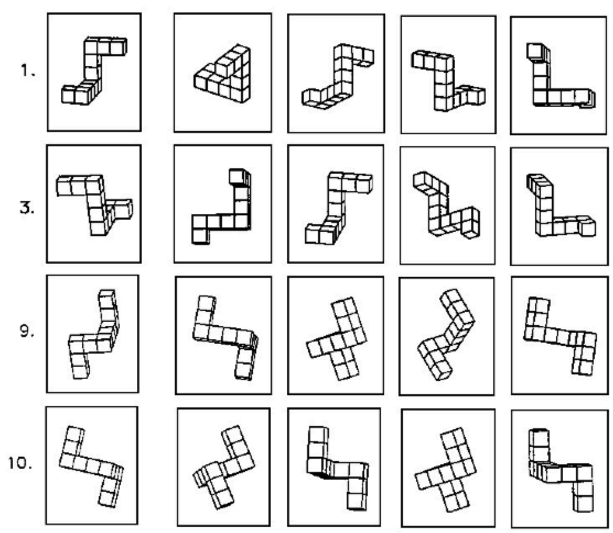

Figure 19. ***Example items from the redrawn Vandenberg & Kuse mental rotation test***

4.4. Results and discussion

4.4.1. Quantitative analysis

We have summarized dependent and independent variables used in this experiment, displayed in Table 5 and Table 6 respectively.

Table 5. Name, description and range of independent variables used in Experiment 2

Variable name	**Description**	**Range**
Rotation condition	Whether or not the participant rotated physically during the trial.	Binary: VISUAL ONLY TURN or REAL TURN
Visual condition	The visual environment presented to the participant during the trial.	Ternary: GRASS, GRASS WITH LINE, CITY
Turning angle	Turning angle of the path travelled during the trial.	Five levels: 10, 50, 90, 130 or 170 degrees.
Turning direction	Turning direction of the path travelled during the trial.	Binary: LEFT or RIGHT.
Order of visual condition	The order in which a participant experienced visual environments during the experiment.	Binary: GRASS → GWL→ CITY or CITY→ GWL→ GRASS.
Order of rotation condition	The order in which a participant experienced rotations during the experiment.	Binary: REAL TURN FIRST or VISUAL ONLY TURN FIRST.

Table 6. Name, description and range of dependent variables used in Experiment 2

Variable name	Description	Range	Data type
Non-Turner behaviour	For each trial, a pointing response was categorized as Turner or Non-Turner behaviour based on whether the participant pointed in the correct hemisphere or not. If the participant responded as a Non-Turner in that trial, this variable was set to 1, otherwise 0.	For each trial: 0 - 1, binary. When averaged: 0 - 1, continuous.	Ratio
Response time	Time passed, in seconds, from the moment a participant was instructed to point until she had settled on a pointing direction and a pointing response was registered.	0.2 - 10 seconds.	Ratio
Task difficulty	Responses from participants when asked to rate the difficulty of the previous trial (only after trials with 50 and 130 degree turns).	0 - 100, continuous.	Ratio
Condition difficulty	Responses from participants when asked to rate the difficulty of each of the six condition combinations (2 rotation conditions X 3 visual conditions) in the post-experimental interview.	0 - 10, integer.	Ratio
Mental rotation test score	Scores of a mental rotation test after being completed successfully by the participant.	0 - 24, integer.	Ratio
Location memory test score	Scores of a location memory test after being completed successfully by the participant.	0 - 27, integer.	Ratio
Visualization ability	Self-rated estimates of participant's own visualization abilities.	0 – 10, integer	Ratio
Spatial orientation ability	Self-rated estimates of participant's own spatial orientation abilities.	0 – 10, integer	Ratio

Response time and *Task difficulty rating* were analyzed using generalized linear mixed models with a split-split-split design structure. *Condition difficulty rating* was analyzed using a generalized mixed model with a split-split design structure.

As *Non-Turner behaviour* is a binary variable, an attempt was made to fit a logistic regression mixed model with a split-split-split design structure. However, the estimation routines for the logistic regression did not converge. Therefore, *Non-Turner behaviour* was fit using an ordinary linear mixed model with a split-split-split design structure. Although not ideal, tests that result from this assumption have in similar situations been shown to be as good as, or better than, tests that result from the traditional analysis (Fang & Loughin, 2013). We believe that the results from this analysis will represent a fair approximation to the truth.

The models for the dependent variables Response time (Table 8), Task difficulty rating (Table 9), and Non-Turner behaviour (Table 7) included the independent variables Order of visual condition, Visual condition, Rotation condition, Turning angle, and all possible interactions as explanatory variables. The models for Condition difficulty rating (Table 10) included Order of visual condition, Visual condition, and Rotation condition and all possible interactions as explanatory variables.

In addition, the variables *Mental rotation test score*, *Location memory test score, Visualization ability* and *Spatial orientation ability* were tested in all models as potential covariates. The final form of each model was determined by performing backward elimination of the covariates as follows:

1. Fit the full model including covariates.
2. Delete the covariate with the largest non-significant *p*-value from the model (alpha-level was set to 0.05).
3. Refit the reduced model.
4. Repeat steps 2 and 3 until all remaining covariates have *p*-values less than 0.05 or have been eliminated from the model.

Using backwards elimination, all covariates were eliminated from the final models for *Response time* and *Condition difficulty rating*. The final models for *Non-Turner behaviour* and *Task difficulty rating* included the covariate *Mental rotation test score*; other covariates were eliminated. In other words, no significant effect was found for *Location memory test score*, *Visualization ability* or *Spatial orientation ability*. The explanatory variables and their interactions were retained in the model regardless of significance. All computations were performed using SAS/PROC MIXED in Version 9.3 of the SAS System (SAS Institute Inc, Cary, NC, USA).

Effects from this model are summarized in Tables 7-10. In the following subsection, we will then discuss results from our inferential statistical analysis in the context of individual hypotheses described in subsection 4.2.

Table 7. F-test results for Dependent Variable Non-Turner Behaviour (significant effects boldfaced)

Effect	Numerator Degrees of Freedom	Denumerator Degrees of Freedom	F Value	p-value
Order of visual condition	1	33	4.54	**0.0406**
Visual Condition	2	68	25.69	**<.0001**
Order of Visual Condition * Visual Condition	2	68	4.59	**0.0135**
Rotation Condition	1	102	0.05	0.8252
Order of Visual Condition * Rotation Condition	1	102	3.14	0.0795
Visual Condition * Rotation Condition	2	102	0.92	0.4009
Order of Visual Condition * Visuals * Rotation	2	102	0.09	0.9095
Turning Angle	4	1836	8.78	**<.0001**
Order of Visual Condition * Turning Angle	4	1836	5.29	**0.0003**
Visual Condition * Turning Angle	8	1836	2.74	**0.0052**
Order of Visual Condition * Visual Condition * Turning Angle	8	1836	3.01	**0.0023**
Rotation Condition * Turning Angle	4	1836	0.49	0.7420
Order of Visual Condition * Rotation * Turning Angle	4	1836	0.33	0.8595
Visual Condition * Rotation Condition * Turning Angle	8	1836	2.28	**0.0201**
Order of Visual Condition * Visual Condition * Rotation Condition * Turning Angle	8	1836	0.72	0.6780
Mental Rotation Score	1	33	8.32	**0.0069**

Table 8. F-test results for Dependent Variable Response Time

Effect	Numerator Degrees of Freedom	Denumerator Degrees of Freedom	F Value	p-value
Order of Visual Condition	1	34	2.69	0.1099
Visual Condition	2	68	6.03	**0.0039**
Order of Visual Condition * Visual Condition	2	68	1.51	0.2285
Rotation Condition	1	102	0.26	0.6084
Order of Visual Condition * Rotation Condition	1	102	0.72	0.3979
Visual Condition * Rotation Condition	2	102	0.93	0.3991
Order of Visual Condition * Visual Condition * Rotation Condition	2	102	2.62	0.0780
Turning Angle	4	1836	2.97	**0.0185**
Order of Visual Condition * Turning Angle	4	1836	0.84	0.5021
Visual Condition * Turning Angle	8	1836	0.52	0.8416
Order of Visual Condition * Visual Condition * Turning Angle	8	1836	1.06	0.3917
Rotation Condition * Turning Angle	4	1836	0.65	0.6246
Order of Visual Condition * Rotation Condition * Turning Angle	4	1836	0.73	0.5743
Visual Condition * Rotation Condition * Turning Angle	8	1836	1.16	0.3170
Order of Visual Condition * Visual Condition * Rotation Condition * Turning Angle	8	1836	0.95	0.4700

Table 9. F-test results for Dependent Variable Task Difficulty

Effect	Numerator Degrees of Freedom	Denumerator Degrees of Freedom	F Value	p-value
Order of Visual Condition	1	33	1.27	0.2678
Visual Condition	2	68	30.03	**<.0001**
Order of Visual Condition * Visual Condition	2	68	2.55	0.0853
Rotation Condition	1	102	0.21	0.6443
Order of Visual Condition * Rotation Condition	1	102	0.04	0.8355
Visual Condition * Rotation Condition	2	102	1.28	0.2825
Order of Visual Condition * Visual Condition * Rotation Condition	2	102	0.00	0.9963
Turning Angle	1	612	16.57	**<.0001**
Order of Visual Condition * Turning Angle	1	612	0.35	0.5557
Visual Condition * Turning Angle	2	612	1.25	0.2869
Order of Visual Condition * Visuals * Turning Angle	2	612	1.22	0.2945
Rotation Condition * Turning Angle	1	612	0.99	0.3197
Order of Visual Condition * Rotation Condition * Turning Angle	1	612	0.01	0.9376
Visual Condition * Rotation Condition * Turning Angle	2	612	0.38	0.6815
Order of Visual Condition * Visual Condition * Rotation Condition * Turning Angle	2	612	1.23	0.2926
Mental Rotation Score	1	33	8.21	**0.0072**

Table 10. ***F-test results for Dependent Variable Condition Difficulty as rated in the post-experimental interview***

Effect	**Numerator Degrees of Freedom**	**Denumerator Degrees of Freedom**	**F Value**	**p-value**
Order of Visual Conditionual Condition	1	34	0.90	0.3500
Visual Condition	2	68	41.28	**<.0001**
Order of Visual Condition*Visual Condition	2	68	0.73	0.4862
Rotation Condition	1	2046	0.18	0.6696
Order of Visual Condition*Rotation Condition	1	2046	198.37	**<.0001**
Visual Condition*Rotation Condition	2	2046	2.23	0.1076
Order of Visual Condition*Visual Condition*Rotation	2	2046	1.78	0.1696

Descriptive analysis of pointing responses

As detailed in section 4.1 we designed Experiment 2 to investigate the spatial updating process through the lens of Turner vs. Non-Turner, by categorising each pointing response – rather than a participant as a whole – as conforming to either a Turner or Non-Turner strategy. Our method of categorisation was quite simple: If a participant pointed in the overall correct hemisphere, that pointing response was categorised as a Turner response. If they pointed in the overall incorrect hemisphere, that pointing response was categorised as a Non-Turner response (see Figure 9 for a visual explanation).

As far as we know this is a novel approach, and so we were interested in reviewing the overall trends in pointing response categorisations for each participant. Did pointing behaviour, as categorised by our method, change enough between trials to justify this method of categorising each trial rather than each participant? A descriptive analysis (see Figure 20) indicates that many participants did indeed exhibit different pointing behaviours between trials, with overall participant Non-Turner rate ranging from 0% to 85% (mean = 19.8%).

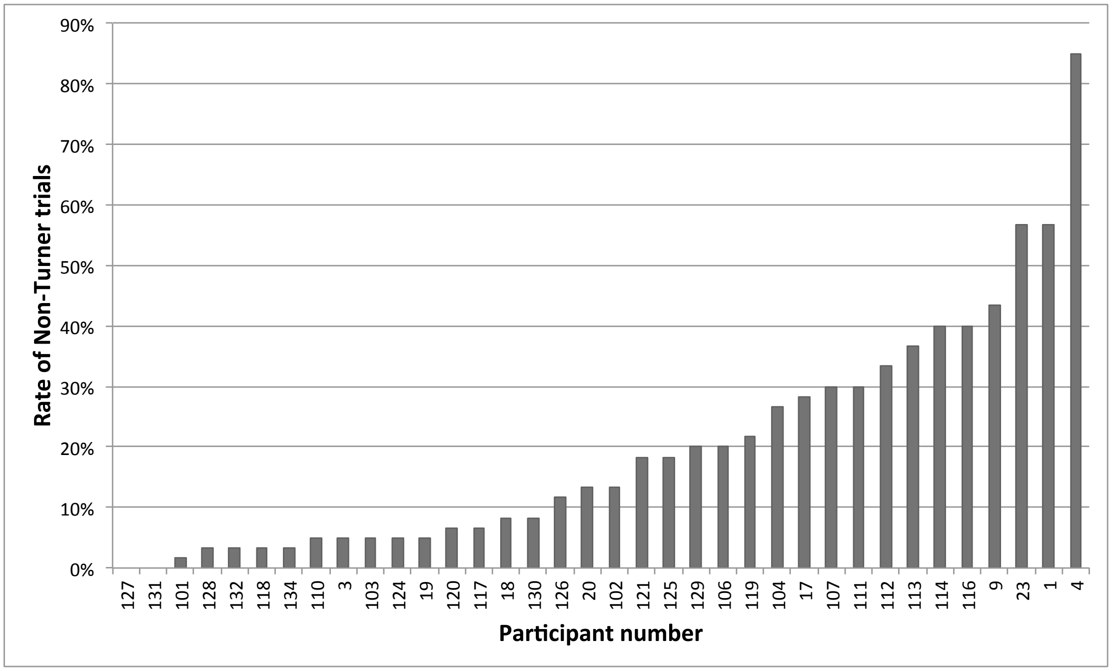

Figure 20.*   *Mean rate of Non-Turner trials for each participant

This would provide validation for our choice of pointing response categorisation, however we must first eliminate the possibility of participants simply pointing incorrectly into the incorrect hemisphere when actually intending to point into the correct one, and vice versa. In practice, this only applies to the smallest and largest turning angles where a relatively small error in pointing could result in an incorrect categorisation. We therefore reviewed the overall trends in pointing response categorisations for each participant, with trials with 10° and 170° are excluded (see Figure 21). Although the mean Non-Turner rate is slightly lower (17.1%), we can see that many participants did indeed exhibit quite mixed pointing behaviour across trials.

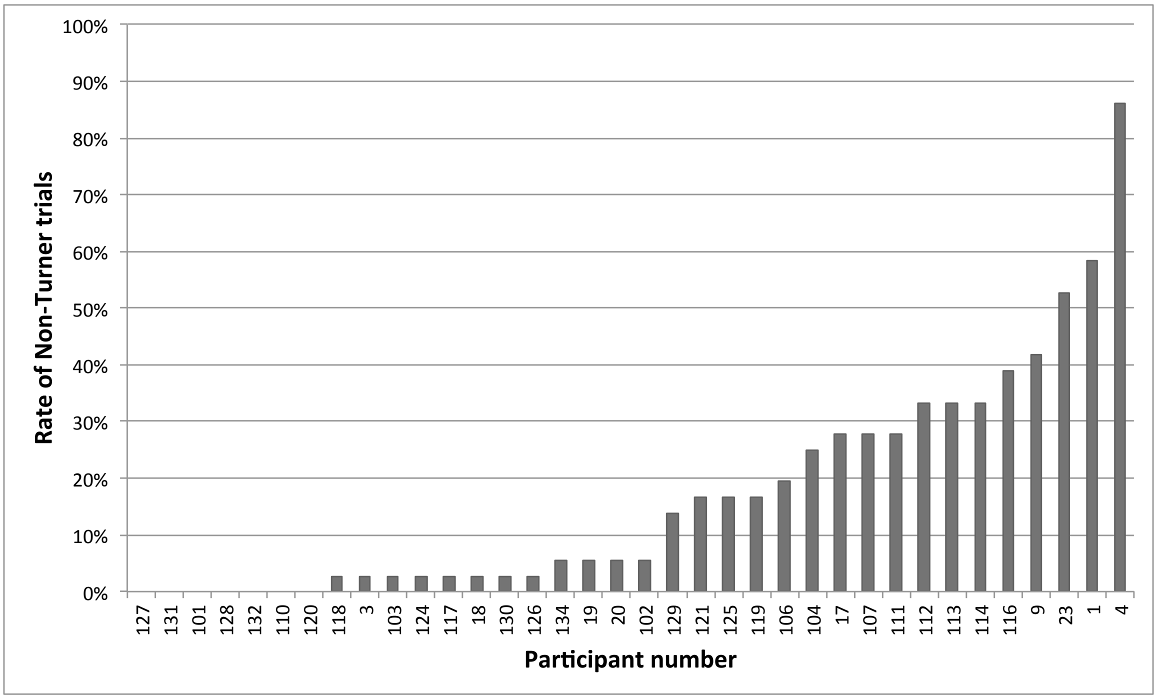

Figure 21. Mean rate of Non-Turner trials for each participant (10° and 170° turns excluded)

Finally, we were interested in the absolute pointing error for each participant – the metric used to estimate spatial orientation in Experiment 1 – so that we could compare it with the novel categorisation approach – rate of Non-Turner trials – used in Experiment 2. As both variables are intended to measure the same thing, we should see a correlation between them when plotted against each other. As shown in Figure 22, there is a clear trend where participants with low absolute pointing errors exhibit a low rate of Non-Turner trials, and vice versa (R^2 = .857). We can therefore conclude that our choice, and implementation, of pointing behaviour categorisation is adequate for the inferential analysis described below.

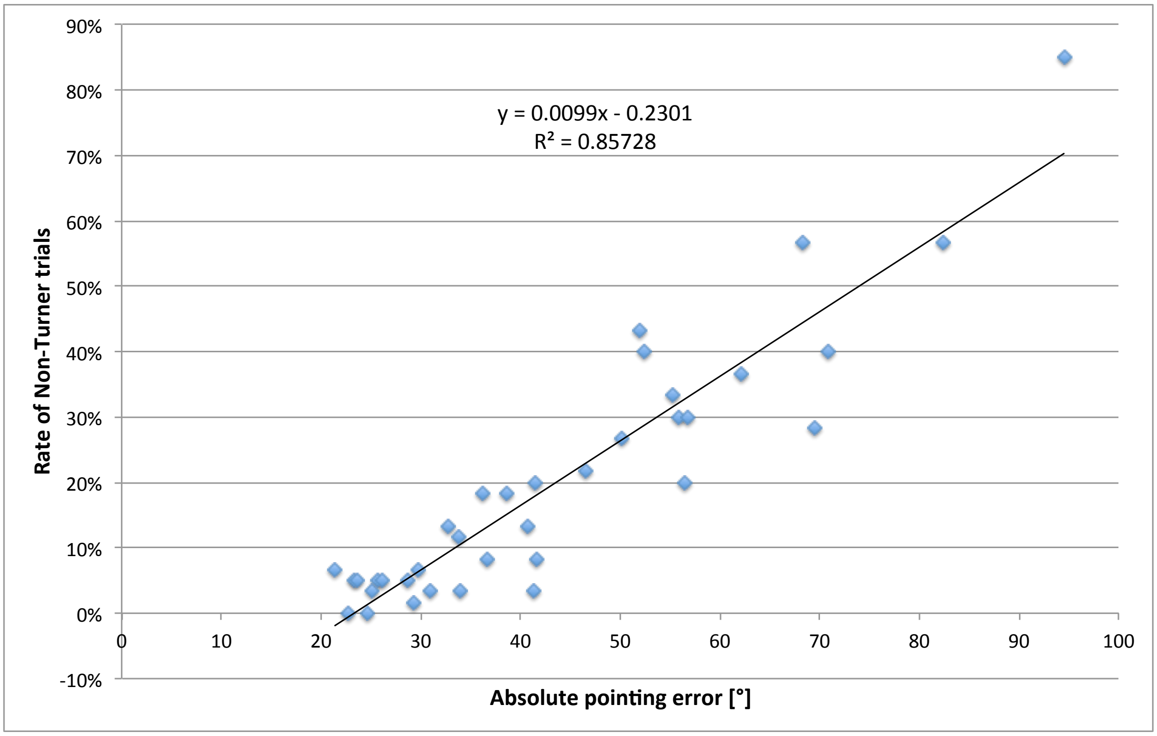

Figure 22. Mean rate of Non-Turner trials plotted against mean absolute pointing error

Note. Each point represents all pointing responses from one participant.

Hypotheses

H1: Added physical rotations increase Turner behaviour

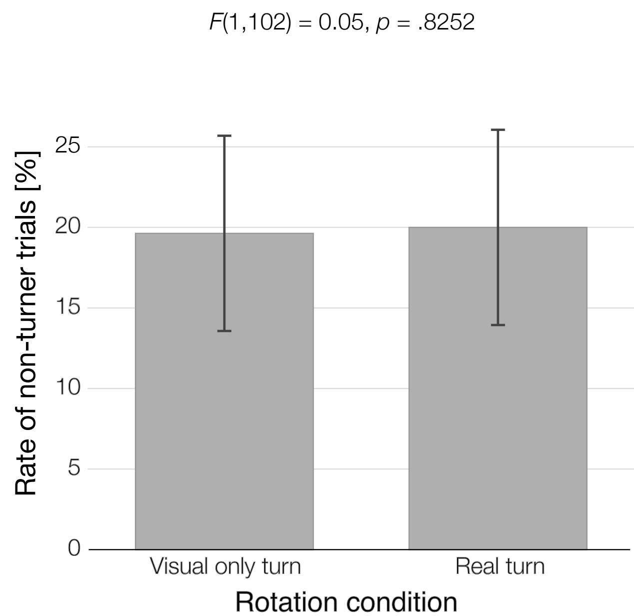

Figure 23. Mean rate of Non-Turner trials for each rotation condition

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot.

We found no significant effect of *Rotation condition* on *Non-Turner behaviour*, $F(1, 102) = .05$, $p = .8252$. That is, participants did not exhibit any more or less Non-Turner behaviour during trials where physical motion cues were available, compared to trials with no physical motion cues (see Figure 23).

These findings contrast with the prevailing notion that physical motion cues are necessary (Klatzky et al., 1998; Ruddle & Lessels, 2006; Wraga et al., 2004) and sufficient (Farrell & Robertson, 1998; Presson & Montello, 1994; Rieser, 1989) to enable spatial updating. Our results therefore support and extend the challenging notion that physical motion cues can be insufficient to enable spatial updating (Riecke et al., 2007; Riecke, Heyde, et al., 2005), even while travelling on a curvilinear path trajectory including rotations and translations (Riecke et al. only used rotations).

H2: Added physical rotations make the task feel easier

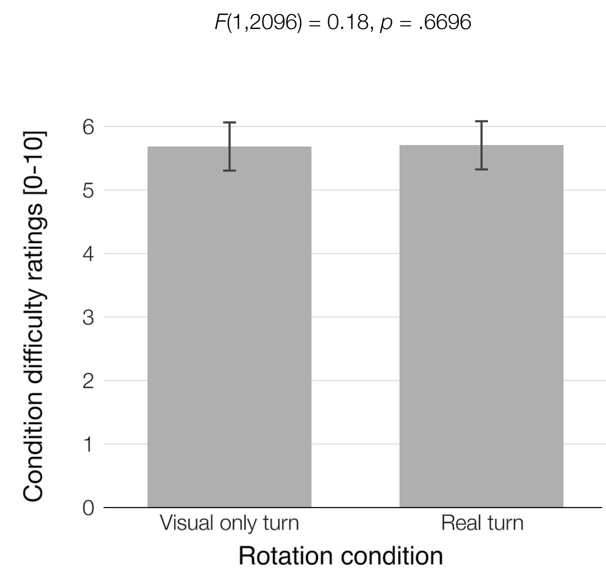

Figure 24. Mean condition difficulty ratings for each rotation condition

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot.

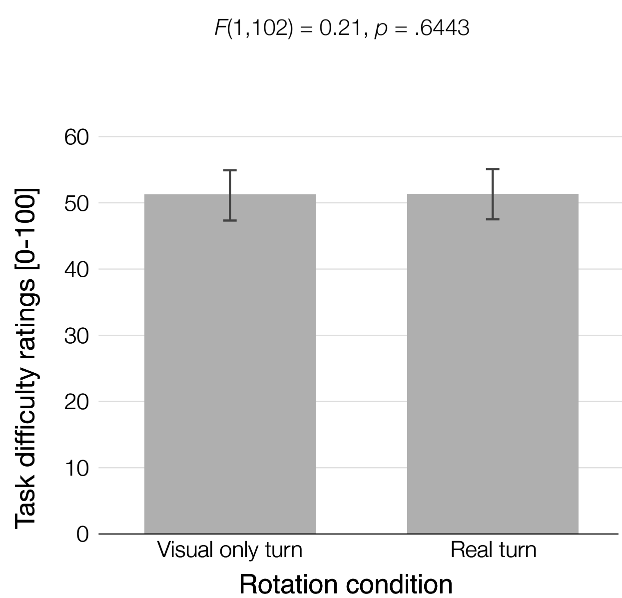

Figure 25. Mean task difficulty ratings for each rotation condition

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot.

We found no significant effect of *Rotation condition* on *Task difficulty rating*, $F(1, 102) = .21$, $p = .644$, or *Condition difficulty rating*, $F(1, 2046) = .18$, $p = .6696$. In other words, participants found conditions with physical rotations neither easier nor harder than conditions with no physical rotations (see Figure 24 and Figure 25).

These findings demonstrate that participants perceived the added information provided by the physical motion cues used in this experiment as neither helpful nor unhelpful. Again, these findings contrast with the prevailing notion that physical motion cues are necessary (Klatzky et al., 1998; Ruddle & Lessels, 2006; Wraga et al., 2004) and sufficient (Farrell & Robertson, 1998; Presson & Montello, 1994; Rieser, 1989) to enable spatial updating, and support the challenging notion that physical motion cues can be insufficient to enable spatial updating (Riecke et al., 2007; Riecke, Heyde, et al., 2005), even while travelling on a curvilinear path trajectory including rotations and translations (Riecke et al. only used rotations). Furthermore, these results corroborate those from our previous experiment; that physical rotations do not seem to matter for this kind of task/environment.

H3. Naturalistic and structured visuals increase Turner behaviour

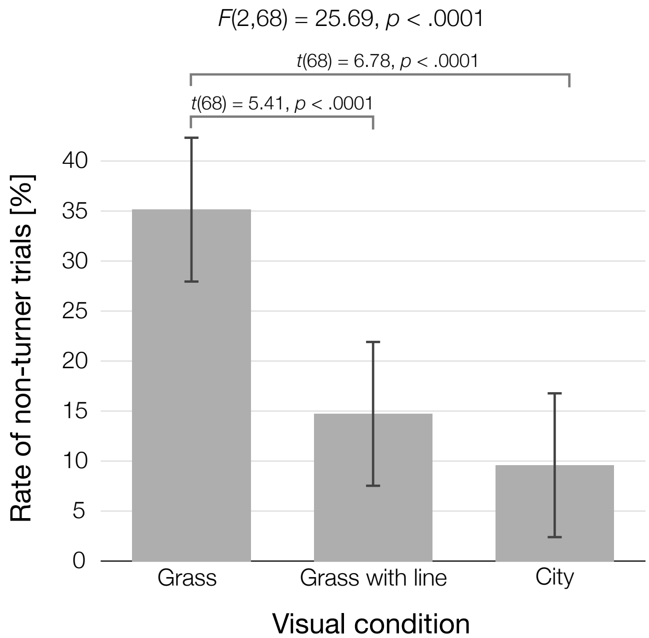

Figure 26. Rate of Non-Turner trials for each visual condition

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot. Significant pairwise differences (with t-test results) are displayed inline for each pair.

Visual condition significantly affected the *Non-Turner behaviour* variable, $F(2, 68) = 25.69$, $p < .0001$ (see Figure 26). Tukey-Kramer adjusted post hoc tests showed that Non-Turner rates were 11.3 percentage points higher for the GRASS environment (LS mean: 35.14%, SE: 3.61%) compared to the GRASS WITH LINE environment (LS mean: 14.72%, SE: 3.61%). Additionally, Non-Turner rates were 16.5 percentage points higher for the GRASS environment (LS mean: 35.14%, SE: 3.61%) compared to the CITY environment (LS mean: 9.58%, SE: 3.61%). All reported differences are at the lower 95% confidence limit, i.e. the difference between the lower confidence limit for one variable and the higher confidence limit for the other.

In other words, participants exhibited significantly less Non-Turner behaviour in the GRASS WITH LINE and CITY environments, compared to the GRASS environment. This matches our prediction that naturalistic and structured visuals decrease Non-Turner behaviour, and supports the notion that a structured visual stimulus can suffice to trigger automatic spatial updating even without physical motion cues.

However, we did not predict that the GRASS WITH LINE environment would be in such stark difference with the GRASS environment, given that only difference between the two is a thin white line that traces the next few meters of the path and gives them advance notice about the direction (but not extent/amount) of the upcoming turn. The fact that this line sufficed to lower the rate of Non-Turner behaviour by more than 11 percentage points suggests that the amount of structure in a visual stimulus may not need to be complex to have the desired effect.

To the best of our knowledge, this has not been reported before in the literature. In a similar point-to-origin experiment, Riecke (2008) had a condition where participants received explicit advance information about the upcoming turn, including its direction and magnitude (e.g. “120° left”, delivered verbally), resulting in a significant decrease in pointing errors and pointing variability. In this study however, participants only received advance information regarding the direction of the upcoming turn, but not the magnitude.

H4: Naturalistic and structured visuals decrease task difficulty

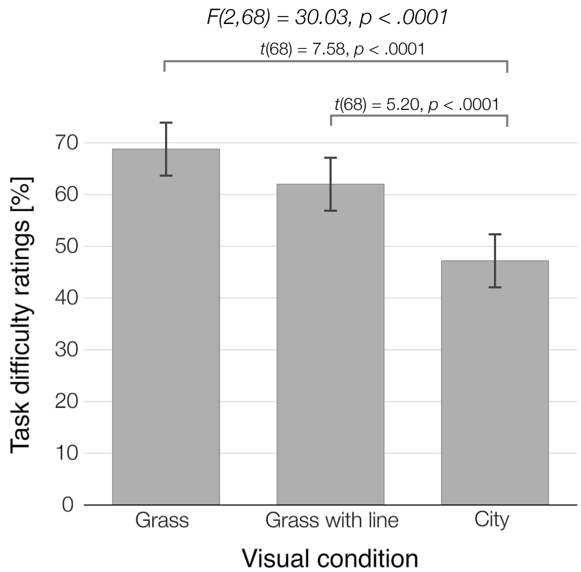

Figure 27. Task difficulty ratings for each visual condition

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot. Significant pairwise differences (with t-test results) are displayed inline for each pair.

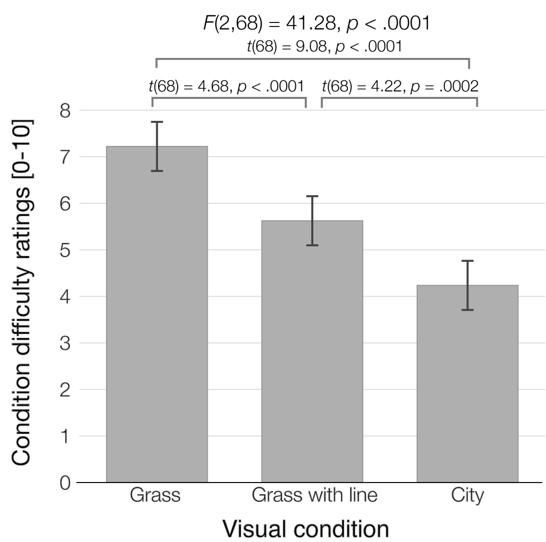

Figure 28. Condition difficulty ratings for each visual condition

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot. Significant pairwise differences (with t-test results) are displayed inline for each pair.

Visual condition significantly affected the *Task difficulty* ratings collected after each of the 50 and 130 degree turns, $F(2, 68) = 30.03$, $p < .0001$ (see Figure 27).

Tukey-Kramer adjusted post hoc tests showed that *Task difficulty rating* was 14.8 percentage points higher for the GRASS environment (LS mean: 68.78%, SE: 2.57%) compared to the CITY environment (LS mean: 47.19%, SE: 2.57%), and 8.0 percentage points higher for the Grass with Line environment (LS mean: 62.00%, SE: 2.57%) compared to the City environment (LS mean: 47.19%, SE: 2.57%). All reported differences are at the lower 95% confidence limit, i.e. the difference between the lower confidence limit for one variable and the higher confidence limit for the other.

Additionally, *Visual condition* significantly affected the *Condition difficulty ratings* collected after the experiment, $F(2, 68) = 41.28$, $p < .0001$, (see Figure 28). Tukey-Kramer adjusted post hoc tests showed that *Condition difficulty rating* was 22 percentage points higher for the GRASS environment (LS mean: 7.22, SE: 0.27 on a 0-10 scale), compared to the CITY environment (LS mean: 4.24, SE: 0.27); 8.1 percentage points higher for the GRASS environment (LS mean: 7.22, SE: 0.27) compared to the GRASS WITH LINE environment (LS mean: 5.63, SE: 0.27); and 6 percentage points higher for the GRASS WITH LINE environment (LS mean: 5.63, SE: 0.27) compared to the CITY environment (LS mean: 4.24, SE: 0.27). All reported differences are at the lower 95% confidence limit, i.e. the difference between the lower confidence limit for one variable and the higher confidence limit for the other.

In other words, participants found the CITY environment significantly less difficult than the other environments, whether they were asked to rate each trial during the experiment (task difficulty ratings) or the visual environment as a whole after the experiment (condition difficulty ratings). Furthermore, they found the GRASS WITH LINE environment less difficult than GRASS when asked to rate the visual environment as a whole after the experiment. This matches our prediction that naturalistic and structured visuals decrease difficulty, and lends further support to the notion that a structured visual stimulus can suffice to trigger automatic spatial updating even without physical motion cues (Riecke et al., 2007; Riecke, Heyde, et al., 2005).

H5: Naturalistic and structured visuals decrease response times

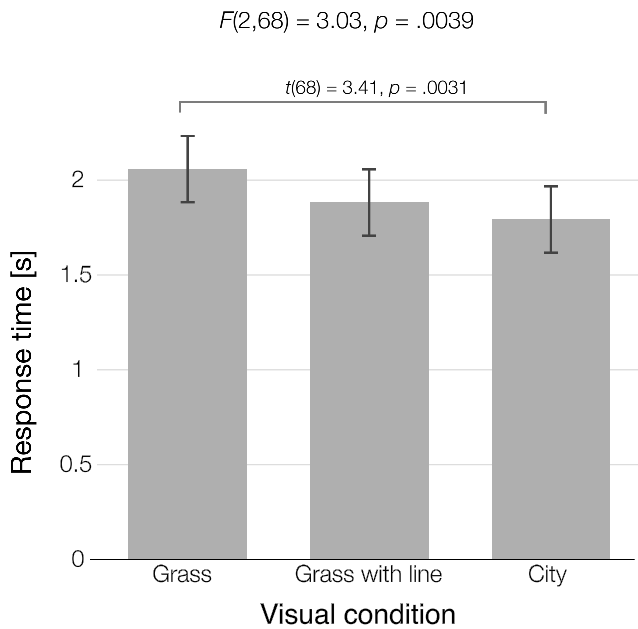

Figure 29. ***Response time for each visual condition***

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot. Significant pairwise differences (with t-test results) are displayed inline for each pair.

Visual condition significantly affected *Response time*, *F*(2, 68) = 3.03, *p* = .0039, (see Figure 29). Tukey-Kramer adjusted post hoc tests showed that *Response time* was 0.11 seconds faster for the City environment (LS mean: 1.79, SE: .09) compared to the Grass environment (LS mean: 2.06, SE: .09). All reported differences are at the lower 95% confidence limit, i.e. the difference between the lower confidence limit for one variable and the higher confidence limit for the other.

While participants pointed slightly slower overall in Experiment 2 (mean: 1.91 seconds) compared to Experiment 1 (mean: 1.67 seconds), this difference can be attributed to the different pointing methods used: In Experiment 1, a pointing response was registered once a modified joystick had been deflected beyond a certain point, while in Experiment 2 a pointing response was registered once the hand-held pointing stick had been deflected beyond a certain point and its motion had settled. The overall time delay introduced by the settling mechanism was typically .1 - .4 seconds, which coincides well with the fact that participants in Experiment 2 pointed on average .24 seconds slower than participants in Experiment 1.

According to the literature, no additional processing time should be needed at the end of the motion used in this experiment – if spatial updating is automatic – as the participant's mental representation of the environment was already updated during the motion (Farrell & Robertson, 1998; Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004). Although quite small, the difference in response time reported above indicates that the naturalistic and structured visuals of the CITY environment were more effective in eliciting automatic spatial updating than the sparse optic flow visuals of the GRASS environment. This supports Riecke et al. (2007), who found a similar benefit when comparing naturalistic and structured visuals with optic flow, and extends their results to include different paths (curvilinear) as well as a different task (point-to-origin instead of pointing to previously learned landmarks).

H6: Larger turning angles do not increase response times

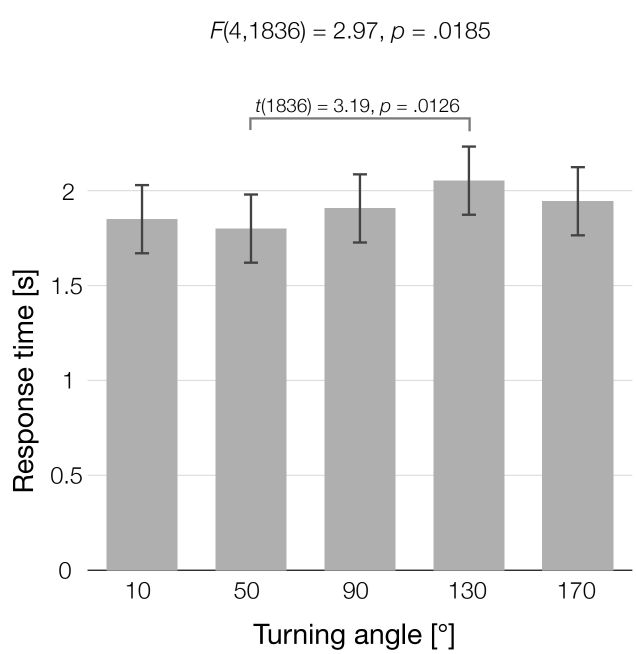

Figure 30. Response time for each turning angle

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot. Significant pairwise differences (with t-test results) are displayed inline for each pair.

Turning angle significantly affected *Response time*, $F(4, 1836) = 2.97$, $p = .0185$, (see Figure 30). Tukey-Kramer adjusted post hoc tests showed that *Response time* was 0.04 seconds faster for 50-degree turns (LS mean: 1.8, SE: 0.09) compared to 130-degree turns (LS mean: 2.05, SE: 0.09). All reported differences are at the lower 95% confidence limit, i.e. the difference between the lower confidence limit for one variable and the higher confidence limit for the other. A linear correlation analysis revealed that *Response Time* increased with larger turning angles overall, $t(215) = 2.43$, $p = .016$, with $\eta^2 = 2.7\%$.

We can therefore conclude that *Response Time* did increase with turning angle, however the correlation explains only 2.7% of the variability of the data. Furthermore, we found no significant difference in *Response Time* for the smallest turn compared to the largest turn. Taken together, these findings suggest that the small increase in *Response Time* is unlikely to be caused by a failure to trigger automatic spatial updating.

H7: Participants overestimate the magnitude of turning angles

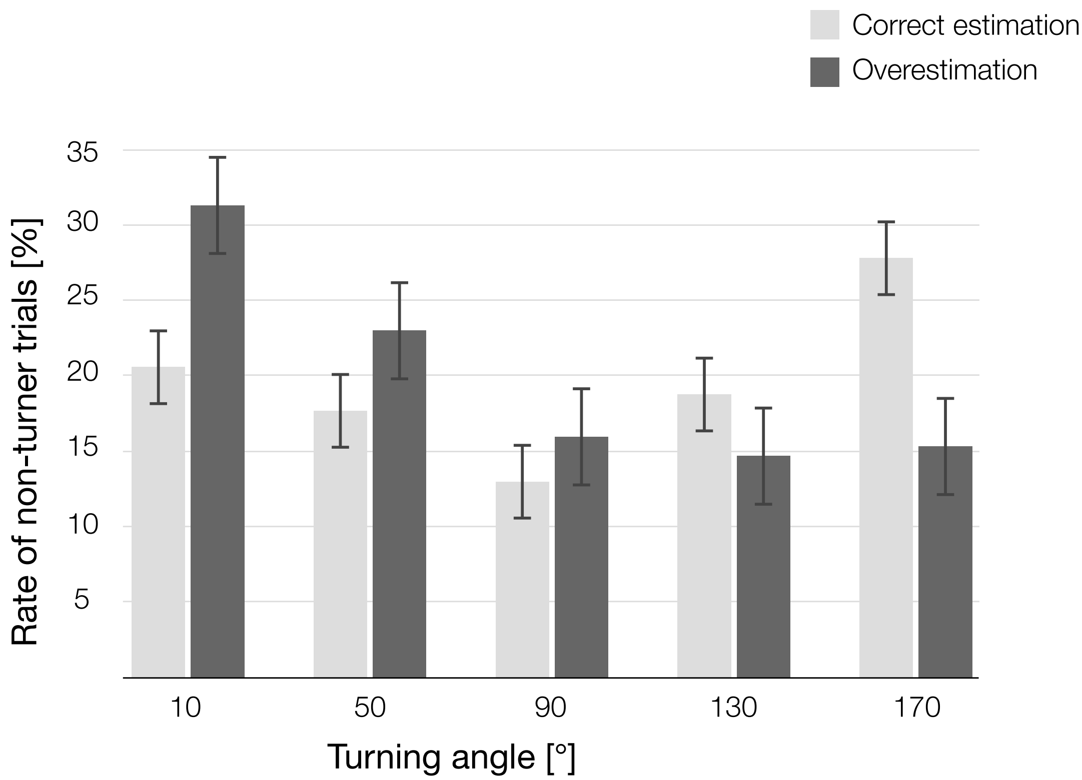

Figure 31. ***Rate of Non-Turner trials for each turning angle, categorized by participant estimation of largest turning angle***

Note. Whiskers indicate one standard error of the mean.

As detailed in subsection 4.3.3, participants were asked to draw their estimation of travelled path trajectories after the experiment. Analysis of each participant's drawing of the largest turning angle revealed that out of 36 participants, 23 perceived the largest turn used in the experiment as over 180 degrees and were categorized as overestimating the largest turn. For this group, the average estimated largest turn was 282 degrees. The 13 participants who did not overestimate the largest turn had an average estimation of 167 degrees. Overall, estimation ranged from 110-360° (M = 221°), that is, participants overestimated the largest turning angle by 30% on average and 111% maximally. These results are comparable to previous studies, such as Riecke (2008) who measured post-experimental turn overestimation at 66% on average 165% maximally.

By way of descriptive analysis (see Figure 31), we can see a clear difference between the Non-Turner vs. Turner pointing behaviour of participants categorized as overestimators compared to others: Overestimators seem to exhibit more Non-Turner behaviour in trials with smaller turns, but less Non-Turner behaviour in trials with larger turns. This contrast is especially pronounced in the lower and upper extremes (10 and 170 degrees, respectively).

These findings confirm that some participants do indeed overestimate the amount of turning in a VR point-to-origin task such as the one used in this experiment, and that a simple analysis of pencil drawings can suffice to make this distinction. This also suggests that some participants exhibiting Non-Turner behaviour may in fact have correctly updated their heading during the trial, but pointed in the wrong hemisphere because they simply felt they had turned much further than they actually did.

As most studies investigating Turner vs. Non-Turner behaviour do not assess this specifically (a notable exception is Riecke (2008), who found participants to frequently misestimate turning angles in a post-experimental debriefing), it is possible that some of the Non-Turner behaviour in prior studies are a result of this confound: that the overestimation of turns can lead to hemisphere errors which could be misinterpreted as Non-Turner behaviour, even though the participant might actually be exhibiting Turner behaviour but overestimated the turning angle. In order to make a clear and robust categorization of Turner vs. Non-Turner behaviour, the potential misperception of turning angles (clearly shown here and in prior work (Riecke, Schulte-Pelkum, & Bülthoff, 2005; Riecke, 2008), must be treated as a confounding variable and controlled as such. We propose that this method of categorization would be a useful addition to future studies.

H8: Participants "keep" the pointing strategies they adopted in the visual environment they experienced first

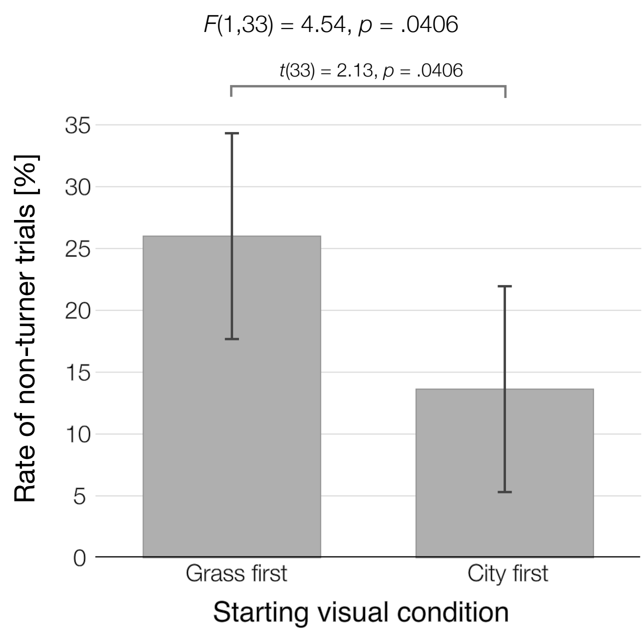

Figure 32. ***Mean rate of Non-Turner trials for each visual starting condition***

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot. Significant pairwise differences (with t-test results) are displayed inline for each pair.

Order of visual condition significantly affected *Non-Turner behaviour* $F(1, 33) = 4.54$, $p < .0406$, (see Figure 32). Tukey-Kramer adjusted post hoc tests showed that participants who started in CITY exhibited significantly less Non-Turner behaviour (LS mean: 13.63%, SE: 4.09%) than participants who started in GRASS (LS mean: 26.0%, SE: 4.09%), with a difference of 0.56 percentage points. All reported differences are at the lower 95% confidence limit, i.e. the difference between the lower confidence limit for one variable and the higher confidence limit for the other.

Although significant, the difference was found to be very small. However, these results indicate that participants who started the experiment in a naturalistic and structured visual environment (CITY) transferred their spatial strategy, in part, to the sparse optic flow visual environments (GRASS WITH LINE and GRASS). See section 5.2 for a detailed discussion on these findings.

H9: Perceived difficulty of rotation conditions depends on the initially experienced visual environment

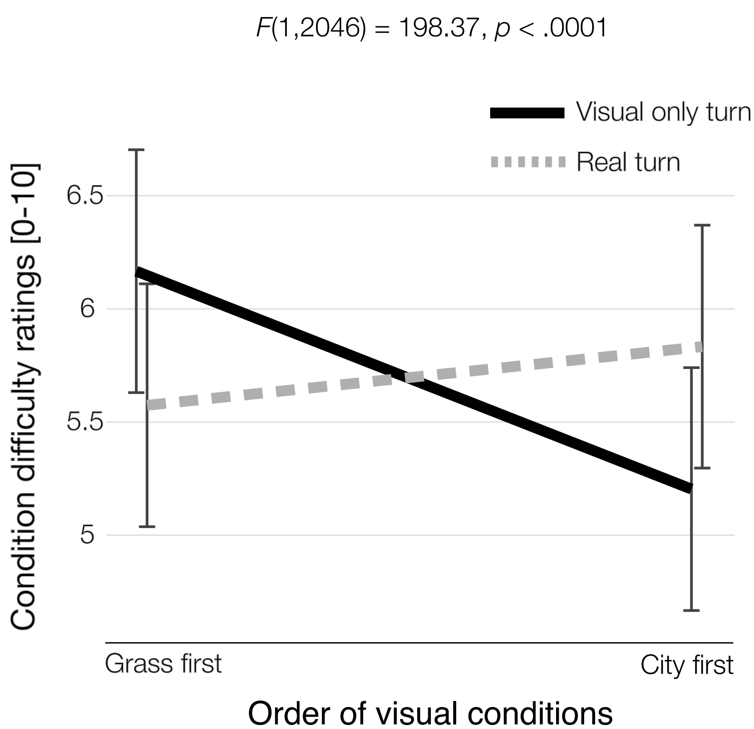

Figure 33. ***Condition difficulty ratings of rotation conditions for different orders of visual conditions***

Note. Whiskers indicate one standard error of the mean. *F* and *p* values are displayed at the top of the plot.

As detailed in our description of hypothesis 8, should we detect a transfer effect of different visual environments, we can measure the contribution of physical motion cues in the participant's mind by looking at how difficult they found conditions with physical motion cues in successive visual environments (see H8 in section 4.2 for details).

We found a significant interaction effect between *Order of visual conditions* and *Rotation condition* on *Condition difficulty rating*, $F(1, 2046) = 198.37$, $p < .0001$, (see Figure 33). Tukey-Kramer adjusted post hoc tests showed that participants who started in CITY found the REAL TURN rotation condition significantly more difficult (LS mean: 5.83, SE: 0.26) than the VISUAL ONLY TURN rotation condition (LS mean: 5.20, SE: 0.26) with a difference of 0.47. Conversely, participants who started in GRASS found the VISUAL ONLY TURN rotation condition significantly more difficult (LS mean: 6.17, SE: 0.26) than the REAL TURN rotation condition (LS mean: 5.57, SE: 0.26) with a difference of 0.43. All reported differences are at the lower 95% confidence limit, i.e. the difference between the lower confidence limit for one variable and the higher confidence limit for the other.

These findings show that the benefit of a physical motion stimulus, in the participant's mind, is partly dependent on the visual stimulus they experienced first. When the initial visual cues are sparse (i.e. GRASS), physical motion cues are perceived as more helpful (in the sense that participants rate conditions with such cues as easier than conditions without), whereas if the initial visual cues are rich (i.e. CITY), physical motion cues are perceived as less helpful.

As our experimental conditions were designed to represent different levels of information density (i.e. the visual information in GRASS is sparse while CITY is rich, and vestibular/proprioceptive information in VISUAL ONLY TURN is sparse while REAL TURN is rich), it is possible that participants starting with sparse visual cues learned to rely more on the vestibular/proprioceptive cues, as they provided information that was absent or difficult to interpret from the visual stimulus. Conversely, those who started with rich visual cues may have learned to ignore, suppress or note take into account the vestibular/proprioceptive cues, as they provided little information that was not already available in the visual stimulus. To the best of our knowledge, this has not been hypothesized before in the literature.

H10: Mental rotation ability correlates with lower Non-Turner rates

Expectedly, we found that *Mental rotation test score* significantly affected *Non-Turner behaviour*, slope b = -.016, $t(33) = -2.88$, $p = .0069$, and *Task difficulty rating*, slope b = -1.13, $t(33) = -2.87$, $p = .0072$. In other words, participants who scored higher on the redrawn Vandenberg & Kuse mental rotation test (Peters et al., 1995) exhibited significantly lower Non-Turner rates and found individual trials easier compared to those who scored lower on the mental rotation test. These results corroborate prior studies showing a positive correlation between mental rotation ability and spatial task performance (Gardony et al., 2014; Malinowski, 2001), providing validation for our method of measuring spatial orientation.

4.4.2. Qualitative analysis

A post-experimental interview was conducted for each participant, where they were asked several questions about their experience during the study to get further insights into potential underlying strategies, relative perceived importance of the different cues, and general feedback. The questions are detailed in Table 11.

Table 11. Post-experimental interview questions

1. How did you solve the task?
2. Did you use any strategies?
3. What did you think of the physical motions used in the study?
4. Did you use the physical motions for spatial orientation?
5. What did you think of the visuals used in the study?
6. Did you use the visuals for spatial orientation?
7. Did anything bother you during the experiment?

These questions were treated as open-ended, and participants were free to speak their mind when responding to each. We then performed a thematic analysis, following the steps outlined in Braun & Clarke (2006), detailed in Table 12.

Table 12. Phases of thematic analysis

Phase	Description of the process
1. Familiarizing yourself with your data	Transcribing data (if necessary), reading and re-reading the data, noting down initial ideas.
2. Generating initial codes	Coding interesting features of the data in a systematic fashion across the entire data set, collating data relevant to each code.
3. Searching for themes	Collating codes into potential themes, gathering all data relevant to each potential theme.
4. Reviewing themes	Checking if the themes work in relation to the coded extracts (Level 1) and the entire data set (Level 2), generating a thematic 'map' of the analysis.
5. Defining and naming themes	Ongoing analysis to refine the specifics of each theme.
6. Producing the report	The final opportunity for analysis. Selection of vivid, compelling extract examples, final analysis of selected extracts, relating back of the analysis to the research question and literature, producing a scholarly report.

After collecting all the participants' responses to these questions, we coded each answer to capture its focal point. If an answer referred to two or more different stimuli or tasks, we coded them so that each coded response pertained to a single topic. For example, the answer "The grass environment was really boring, and the graphics weren't very nice in the city" became two coded responses:

1. "Grass environment was really boring"

2. "Graphics weren't very nice in the city"

Additionally, since participants were allowed to speak their mind freely when answering questions, including revisiting previous questions or discussing related topics, we opted for an inductive thematic analysis approach (e.g. Frith & Gleeson (2004)) where coded responses were treated as an individual comment or feeling regardless of what question prompted the response.

Figure 34. Organization of coded responses into themes

After all responses had been coded and analyzed, we moved to the next step of the thematic analysis: organizing and arranging the responses to discover themes. Using post-it notes and a method of organizing and grouping large amounts of data, called affinity diagram (see Figure 34), we found that responses could be placed in one of two distinct categories: Feelings towards stimuli, or a description of pointing behaviour.

Feelings towards stimuli

Further analysis of responses that described feelings towards a specific stimulus revealed two sub-categories: *feelings of empowerment* and *feelings of impediment*.

Feelings of empowerment

In this category, participants described feelings where they felt more at ease, more comfortable, better able to finish the task, better oriented and so on. These positive feelings described how the relevant stimulus affected the participant in an empowering way. In this subcategory, we identified specific themes, shown in Table 13.

Table 13. Themes of empowerment

Theme	Description	Example
Useful	Answers describing a stimulus as useful for the purposes of finishing the task.	“Rotations were helpful”
Intuitive	Answers describing a stimulus as feeling natural and easy to understand.	“Physical motions felt intuitive”
Immersive	Answers describing a stimulus as giving a sense of “being there”.	“The city gave a feel for the streets”
Informative	Answers describing a stimulus as providing important information.	“The line acted as a sidewalk, giving information”
Orientative	Answers describing a stimulus as improving sense of orientation.	“The line helped me stay oriented”
Belonging	Answers describing a stimulus as feeling familiar, giving a sense of belonging.	“The city was easier because it felt more familiar”

Feelings of impediment

In this category, participants describe feeling uneasy, uncomfortable, less able to finish the task, disoriented, confused and so on. These were negative feelings describing how the relevant stimulus affected the participant in an impeding way. In this subcategory, we identified specific themes, shown in Table 14.

Table 14. Themes of impediment

Theme	Description	Example
Disorienting	Answers describing a stimulus as impairing sense of orientation.	“I felt nauseous in grass and grass with line”
Distracting	Answers describing a stimulus as feeling distracting or misleading.	“I tried to ignore the physical motions”
Confusing	Answers describing a stimulus as confusing or difficult to understand.	“The grass was confusing because of the lack of reference frame”
Irritating	Answers describing a stimulus as annoying.	“The disappearing line was aggravating”
Boring	Answers describing a stimulus as boring or uninteresting.	“The grass was boring”

Having uncovered these recurring themes in participants' responses to how they felt towards stimuli used in the experiment, we evaluated the occurrence of different themes for each stimulus:

Physical rotations

Participants' descriptions of physical rotations was associated with empowering themes of *Useful*, *Immersive* and *Intuitive* (see Figure 35), while many also described impeding themes of *Confusing, Disorienting* and *Distracting* (see Figure 36). Furthermore, no participants described the physical rotations as *Informative*, *Orienting* or *Belonging*. Although almost all participants opined on the physical rotation stimulus, we found no clear trend in their attitudes towards positive or negative feelings, indicating that their feelings on physical motion cues were both mixed and quite polarized. Careful review of participants who voiced positive or negative feelings on the physical rotations also revealed no obvious correlation with starting conditions, either visual or vestibular.

These mixed but strong feelings on physical rotations are somewhat in accordance with the literature, as studies on motion sickness (and the closely related simulator sickness, which does not require physical motion) have shown (Stanney, Mourant, & Kennedy, 1998). We did not anticipate such a clear division however, as the tendency towards negative feelings towards physical motion cues was almost exactly as prevalent as the tendency towards positive feelings.

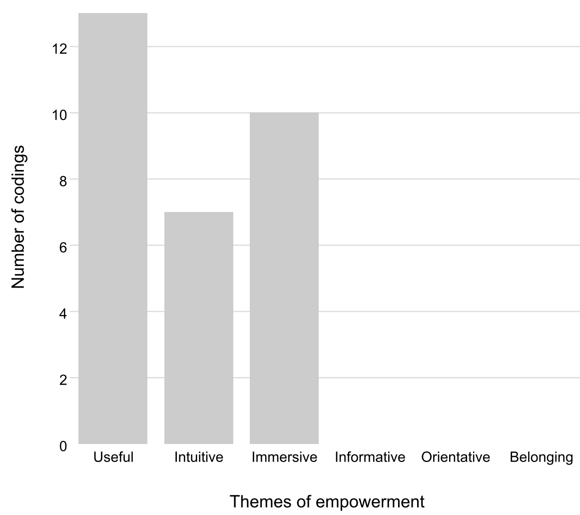

Figure 35. Occurrence of empowering themes in describing physical rotations

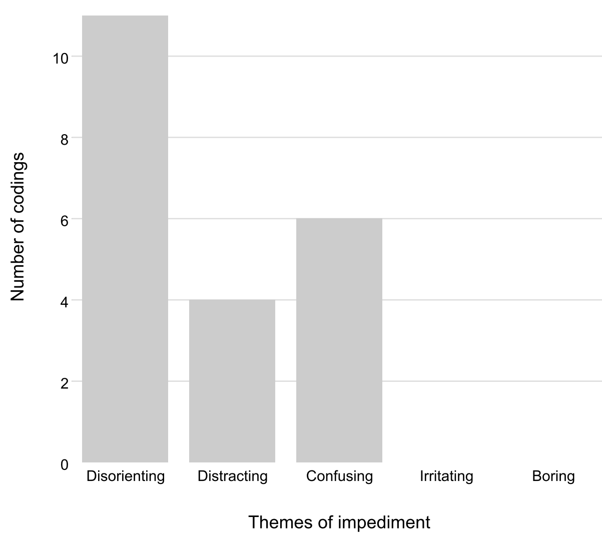

Figure 36. Occurrence of impeding themes in describing physical rotations

Visual environments

Starting with the visual environment of lowest fidelity – GRASS – a single participant noted that they "preferred the grass for its simplicity", but otherwise no feelings of empowerment were associated with the GRASS environment. However, many described feelings of impediment related to this environment, which we associated with the themes of *Disoriented*, *Confused*, *Irritating* and *Boring*. The occurrence of each theme in relation to visual stimuli is depicted in Figure 37 and Figure 38.

For the GRASS WITH LINE environment, many participants' positive responses were associated with the themes *Informative*, *Orientative* and *Useful*, while some reported negative feelings associated with the themes *Boring*, *Confusing* and *Irritating*.

Finally, analysis of feelings towards the CITY environment revealed a strong sense of the themes *Belonging*, *Immersive*, *Informative* and *Useful*. No participants reported feelings of impediment towards the CITY environment.

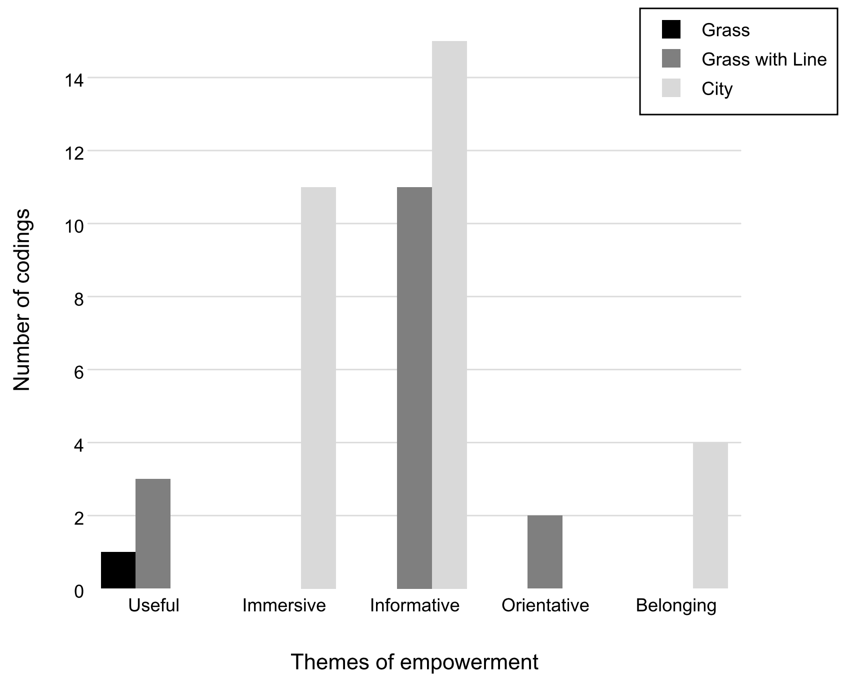

Figure 37. Occurrence of empowering themes in describing visual stimuli

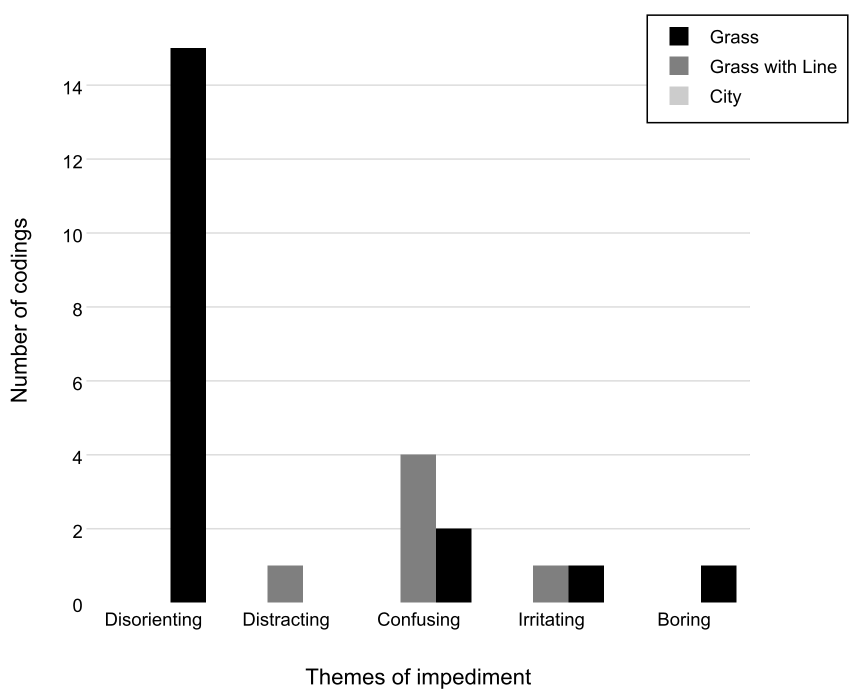

Figure 38. Occurrence of impeding themes in describing visual stimuli

This coincides well with the intended effects of the visuals, where we expected participants to have more difficulty completing the task in the lower-fidelity visual environment while the higher-fidelity visuals would make the task easier. The optic flow environment was designed to give participants the minimum amount of information needed to perceive visual motion, and it is to be expected that in such an environment they might feel confused, disoriented, irritated, bored or otherwise frustrated with the lack of information given to them that they might expect from navigating in most real environments.

Conversely, the CITY environment was designed to give participants the maximum amount of visual information without using landmarks, and so it does not come as a surprise that participants found this amount of information to feel informed and immersed in the environment while completing the task. The reported feelings of belonging were more unexpected, but still fit inside the general idea that the CITY would feel easier and therefore incite more positive feelings.

The fact that there were almost no feelings of empowerment reported for GRASS, and no feelings of impediment reported for CITY make this distinction between the two extremes in the minds of participants even clearer, and coincide with our expectations.

The in-between visual environment – GRASS WITH LINE – was designed as a middle ground between the two extremes, GRASS and CITY, and thus our predictions were less clear. However, the fact that reported feelings for this environment decidedly stemmed from both empowerment and impediment indicates that we found this middle ground at least in participants' minds. Where there was a clear consensus for GRASS and CITY, there was none at all for GRASS WITH LINE; some found it empowering, others impeding.

Taken together, we find this collection of participants' feelings towards the different visual stimuli to show a clear trend of positivity towards higher-fidelity visuals. This matches our quantitative findings quite well, as participants pointed more accurately in higher-fidelity visual environments, and rated them as easier than lower-fidelity ones.

Pointing behaviour

Turning to the other main category – pointing behaviour as reported by participants – we identified some overarching themes that we further analyzed and categorized.

These responses were typically answers to questions 4 and 6 ("Did you use the physical motions for spatial orientation?" and "Did you use the visuals for spatial orientation?"), but as explained above we interpreted each response individually and not just as an answer to a specific question. Examples of these responses are "Looked behind me in city to see the starting point" and "Used different strategies depending on the visuals".

Once we had collected all the comments describing these pointing strategies, we immediately identified three main themes (see Figure 39):

Embodied strategies*:* Strategies where participants focused on something physical, either through external stimulus (e.g. used shoulder to keep track of starting point, focused on physical rotations when available, leaned into rotations, etc.) or internal "gut feeling" (e.g. solved task by feeling, tried to rely on spatial intuition, etc.).

Visual strategies*:* Strategies where participants focused on something visual, either through external stimulus (looking behind them, focusing on the ground, focusing on the line in GRASS WITH LINE condition, etc.) or internal visualization (imagining planting a flag at the starting point, visualizing starting point in their head, imagining seeing themselves from a top-down perspective, etc.)

Mixed strategies*:* Strategies where participants focused on both visual and physical aspects, either simultaneously during the same trial or alternating between trials without a specific preference for a given stimulus.

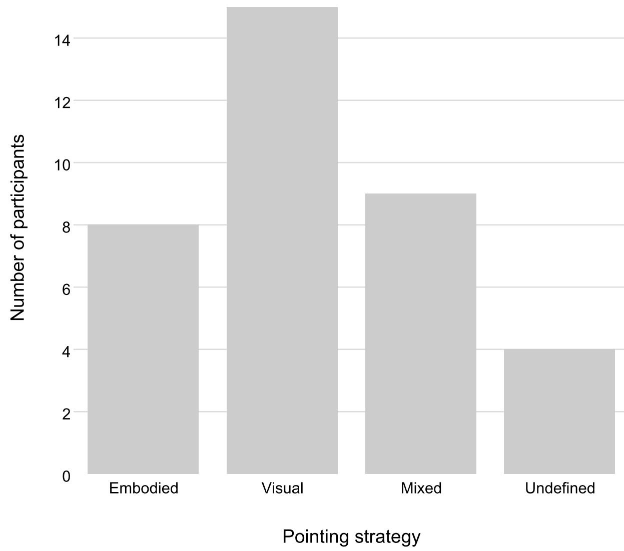

Figure 39. Distribution of participant pointing strategies

Note. Four participants did not describe any pointing strategy, and were categorized as Undefined.

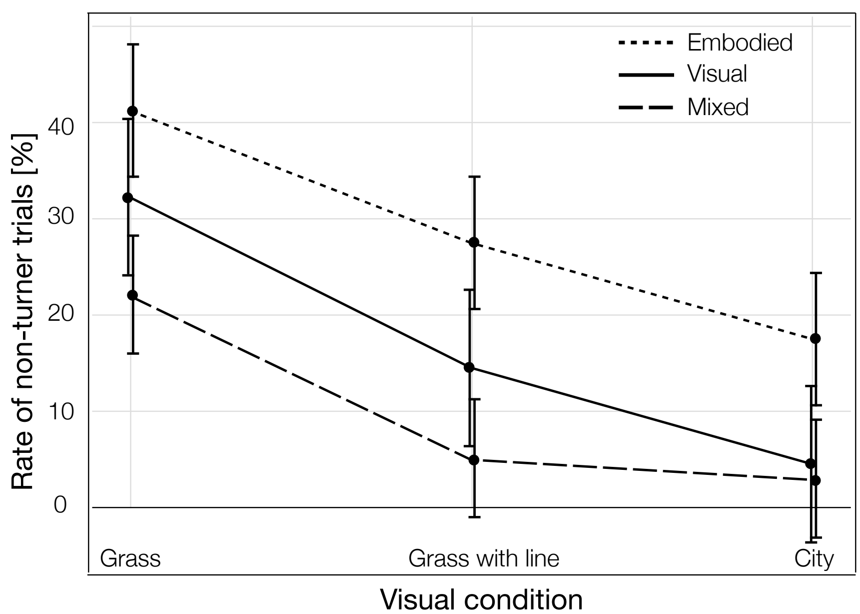

Figure 40. Mean rate of Non-Turner trials for each visual condition, grouped by pointing strategy

Note. Whiskers denote one standard error.

Cross-referencing these coded responses with the Non-Turner dependent variable suggests that participants who mainly described their pointing behaviour in terms of Embodied strategies exhibited Non-Turner behaviour at a higher rate (21%) than participants relying on Visual strategies (15%), with participants describing Mixed strategies exhibiting the lowest Non-Turner rate (9%) (see Figure 40).

Due to the inherent subjectivity of thematic analysis, as well as the groups of unequal sizes, we did not apply inferential statistics on this data. However, this descriptive analysis suggests that individual characteristics in feelings towards stimuli and preferred pointing strategies may significantly impact spatial performance during a point-to-origin task in a virtual environment. Although individual differences such as mental rotation ability, working memory and verbal ability have been found to be a major source of variation in both real-world and computer-mediated spatial tasks (Bryant, 1982; Hegarty, Montello, Richardson, Ishikawa, & Lovelace, 2006; Wolbers & Hegarty, 2010), few studies have investigated these differences in the context of feelings towards a specific stimulus or proclivities for one spatial strategy over another. We propose further studies to explicitly investigate the impact of these characteristics on spatial orientation, so that we can come closer to understanding how our inherent preferences for different spatial strategies combine with external stimuli to keep us oriented (or not) in virtual environments.

5. General discussion and conclusion

This study set out to investigate a long-standing issue for Virtual Reality as a tool for research, design, training, therapy and entertainment: As we move around a virtual environment, we tend to get much more easily lost and disoriented compared to the real world.

In particular, this study focused on the interplay of different stimuli in enabling spatial updating – our innate ability to stay oriented as we move around in the world – specifically the richness and structure of a visual stimulus, and presence of a physical motion stimulus. The prevailing notion in the spatial updating literature is that physical motion cues are sufficient and necessary for triggering automatic spatial updating (Chance et al., 1998; Klatzky et al., 1998; Ruddle & Lessels, 2006; Wraga et al., 2004), but Riecke et al. (2007; 2005) found that visual cues alone can be sufficient to trigger spatial updating provided that they are naturalistic. However, in the studies performed by Riecke et al. participants were only tested with rotational movements, not translations or curved motions. In this study, we extend these results by measuring users' spatial orientation after travelling along curved paths that include translations and rotations.

5.1. Summary of results

In Experiment 1, participants performed a point-to-origin task in a virtual environment, with naturalistic and structured visuals. Measuring participants' spatial orientation via absolute pointing error, we found no effect of adding physical motion cues to match the rotations of the visuals. I.e. the physical motions did not help as far as we could tell, which goes against the prevailing notion that physical motion cues are sufficient and necessary to trigger automatic spatial updating (Chance et al., 1998; Ruddle & Lessels, 2006; Wraga et al., 2004).

Additionally, we observed two participants that exhibited Non-Turner behaviour, that is, they responded as if they were still facing the original direction even though they clearly were aware of the actual path, and continued to exhibit the same behaviour even when additional physical motion cues were provided. To the best of our knowledge, this is the first time that Non-Turner behaviour was reported despite physical motion and naturalistic visuals. This contradicts previous research suggesting that physical motion cues are sufficient to trigger obligatory spatial updating of rotations (e.g., (Farrell & Robertson, 1998)). There are of course differences to prior studies that could have contributed. For example, prior studies that observed automatic spatial updating often alternated translations and rotations (Chance et al., 1998; Farrell & Robertson, 1998; Klatzky et al., 1998; May & Klatzky, 2000; Presson & Montello, 1994; Rieser, 1989; Wraga et al., 2004), whereas we combined them to yield a smoothly curved path that is closer to simulating naturalistic motion.

Our goal with Experiment 2 was to build upon and expand our findings from Experiment 1. Using an upgraded head-mounted-display and pointing device, we had participants do a point-to-origin task in a virtual environment, with differing levels of visual fidelity. In one environment, participants moved around an optic flow environment (GRASS), similar to an infinitely large grass plane with no discerning features or structure. In another environment, they moved around a city similar to the one used in Experiment 1: Structured and naturalistic, yet devoid of landmarks (CITY). In the third environment, participants moved around an optic flow environment similar to the first one, with one key difference: A thin, white line predicted the next few meters of the participant's trajectory, without revealing the entire path and therefore becoming a landmark (GRASS WITH LINE).

In both experiments, we found no effect of physical rotations on spatial orientation, i.e. participants did not exhibit different pointing behaviour during trials where they physically rotated, compared to those with visual-only rotations. This contrasts with the prevailing notion that physical motion cues are sufficient and necessary to enable obligatory spatial updating, and suggests that in some cases they may not suffice. These findings align with the findings of Riecke et al. (2007; 2005) which suggests that physical motion cues may not always be necessary to enable obligatory spatial updating.

Moving our attention to the effects of our visual stimuli, we found that its structure and fidelity had a significant effect on participant's spatial orientation. Participants exhibited significantly less Non-Turner behaviour in the Grass with Line and City environments, compared to the Grass environment. This matched our prediction that naturalistic and structured visuals should decrease Non-Turner behaviour, and supports the notion that a structured visual stimulus can suffice to trigger automatic spatial updating even without physical motion cues.

Finally, a significant contribution of this study is our observation that spatial performance in a visual environment is partly dependent on the environment first experienced by a participant. Those who started the experiment in a sparse optic flow environment (and progressed to a high-fidelity structured environment) exhibited significantly more (almost twice as frequent) Non-Turner behaviour across all conditions than those who started in a high-fidelity environment. Although reported difference at the lower 95% confidence limit was small, we find in this a strong indicator that participants do in some part "keep" a spatial strategy adopted earlier when transitioning to new environments. We also found previously experienced visual stimuli to affect participants' perception of the benefits of physical motion cues. Participants who started in a sparse environment perceived physical motion cues overall to be more helpful than not, whereas those who started in a high-fidelity environment perceived physical motion cues to be less helpful than not.

5.2. Implications and future work

Physical rotations

We found it surprising that physical motion cues did not seem to have any effect at all, even under conditions where the visual stimulus was too sparse to provide useful spatial information. Where Klatzky et al. (1998) found a significant effect of physical motion cues on pointing accuracy in an optic flow spatial orientation task, we found none. However, there are some important differences that potentially explain this discrepancy. First, we must consider the task itself: Path trajectories in Klatzky et al. were two-legged, with a rotation in between, whereas we combined translations and rotations in a smooth curve. Although this combination does make for a trajectory that is closer to real-world travel, it also makes it more difficult to discern between translations and rotations when compared to a trajectory where they have already been separated. As a result, the physical motion cues used in our study may have improved spatial orientation in some cases but worsened it in others. The fact that our qualitative analysis found participants reporting strong feelings – both positive and negative – towards the physical rotations (see section 4.4.2) supports this notion.

Additionally, technical improvements such as vastly increased field of view (44 x 33 degrees in Klatzky et al. vs. 102 x 64 degrees in our study) and a software-controlled rotation platform (compared to manual chair-turning) are likely factors in explaining the difference in these results. We suggest that future studies focus on the parameters of the physical rotations, in order to help us better understand how they affect us. For example, faster rotations might produce different results, as well as slower ones.

The fact that physical motion did not seem to improve performance has implications for the design of virtual reality systems. Our results suggest that expensive motion simulators may not be required for users to navigate effectively in VR, which could drastically lower the financial barrier of entry and open up many possibilities that were either impractical or impossible before. Diminished reliance on large motion-controlled VR setups would also make for a much safer user experience, again offering new use cases that may have been infeasible before. All in all, our findings bode well for designers of virtual environments.

Visuals

We did not predict that the Grass with Line environment would be in such stark difference with the Grass environment, seeing as the only difference between the two is a thin white line that traces the next few meters of the path. The fact that this line alone sufficed to lower the rate of Non-Turner behaviour by more than 11 percentage points suggests that a structured visual environment may not need complex structure to have the desired effect.

How might a thin line in an otherwise sparse optic flow environment be so effective in improving spatial orientation? We cannot conclusively answer this question in this study, however our qualitative findings (see subsection 4.4.2) can provide valuable insight. For example, the Grass with Line environment was strongly associated (i.e. 10 or more codings) with the Informative theme of empowerment, while City was strongly associated with both *Informative* and *Immersive* themes of empowerment. This suggest that participants processed the line on a cognitive level more than a perceptual one, especially when considering that no participants described the line as *Immersive*. Furthermore, the most common impeding theme associated with Grass with Line was *Confusing*, supporting the notion that there is a strong cognitive component to the choice of spatial strategy in the Grass with Line condition.

We propose further research into this phenomenon, for example by varying the manifestation of the information provided by the line. Will a verbal notice elicit a similar effect? Riecke (2008) provided advance verbal notice about the upcoming turn (in a control experimental condition), however they only tested experienced observers (i.e. lab members) and provided information about the direction and magnitude of the turn. Further studies are needed to conclusively answer this question.

On the applied side, these findings indicate that VR designers may benefit from focusing on providing a naturalistic and structured environment, rather than striving for photorealistic shading or other graphical enhancements. If a thin line predicting the upcoming turn direction can improve spatial orientation to such a degree as seen here, what could other manifestations of structured visuals do?

Transfer of strategy between visual environments

Why might the previously experienced visual environment affect our perception of the current one? It is possible that participants starting with sparse visual cues learned to rely more on the vestibular/proprioceptive cues, as they provided information that was absent or difficult to interpret from the visual stimulus. Conversely, those who started with rich visual cues may have learned to ignore or suppress the vestibular/proprioceptive cues, as they provided little information that was not already available in the visual stimulus. To the best of our knowledge, these findings have not been described before in the literature.

Our findings strongly suggest that previous exposure to visual stimuli is a contributing factor to the spatial updating process. We propose further studies to investigate this notion, for example by having participants complete a physical walking task, or even a mental imagination task, before starting a point-to-origin VR experiment similar to the ones described in this study. Finding such tasks to affect spatial performance in the following experiment would suggest that we could bias the selection of spatial strategies through careful selection and adjustment of prior exposure. Another potentially interesting follow-up is to re-analyze previous VR studies in this context.

These results could also prove useful in applied VR uses, such as therapy, entertainment and training (see section 2.1.2 for a detailed overview of practical VR uses). By carefully selecting the initial visual stimulus, designers may be able to subtly steer the user experience in a desired direction, and even compensate for the inherent lack of information afforded by some environments. For example, a flight simulator might use a highly structured environment at the beginning to “prime” trainees, and then advancing them to a more ecologically valid environment such as a featureless sky.

5.3. Limitations

There are limitations to this study, which could have informed or improved our results. First, it is possible that combining visual rotations and translations in a smooth curvilinear path, yet only providing physical rotations, may have been a confounding factor in how the stimuli were perceived by participants. As previous studies have typically split translations and rotations into two separate motions, we must consider this as a potential limitation of this study. On the other hand, a curvilinear path such as used in this study is considerably closer to simulating real-world movement, compared to separated legs of either translations or rotations. One suggestion to future VR researchers is to replicate seminal studies such as performed by Klatzky et al. (1998), with curvilinear excursion paths instead of the two-legged paths used in the original study. This could allow us to better understand how we perceive, combine and process different spatial cues, and move us closer to effective VR navigation.

Another potential limitation of this study is our method of analyzing Turner vs. Non-Turner behaviour on a per-trial basis. Previous studies (Goeke et al., 2013; Gramann et al., 2010, 2005, 2006; Riecke & Wiener, 2007; Riecke, 2008) have typically classified participants as preferring one strategy over the other. We surmised that this per-participant distinction might not as clear-cut as the literature implies, and that other factors than personal preference are at play. In Experiment 2, we therefore chose to categorize each trial, rather than each participant, as conforming to Turner or Non-Turner behaviour, based on whether they pointed in the overall correct hemisphere or not (see “Categorization of pointing responses” under section 4.3.3 for a detailed discussion). This allowed for a finer-grained analysis of how and why people fail to update their rotation.

Our results support this suspicion that the categorization of participants as Turners, Non-Turners or Switchers may be too coarse in some instances. However, due to the novelty of our particular categorization method we cannot rule out potential confounding effects, for example the perception of traveled distance and magnitude of rotations. In fact, we found that many participants did indeed overestimate the magnitude of the largest turn used in the experiment, suggesting that some observed Non-Turner behaviour may have stemmed from this misperception rather than a failure to update headings during rotations.

To advance this line of research, a re-analysis of prior studies using our novel method of per-trial categorization might allow us to better understand why or when people fail to update their rotation. Additionally, analysis of participants' perceived path trajectory such as performed in this study may prove beneficial to tease apart the different factors of overestimation of the traveled path and failure to update headings.

References

Aggarwal, R., Ward, J., Balasundaram, I., Sains, P., Athanasiou, T., & Darzi, A. (2007). Proving the effectiveness of virtual reality simulation for training in laparoscopic surgery. *Annals of Surgery*, *246*(5), 771–9. doi:10.1097/SLA.0b013e3180f61b09

Avraamides, M. N., & Kelly, J. W. (2008). Multiple systems of spatial memory and action. *Cognitive Processing*, *9*(2), 93–106. doi:10.1007/s10339-007-0188-5

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 & Virtual Environments*, *8*(1), 36–53. doi:10.1162/105474699566035

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*, 77–101. doi:10.1191/1478088706qp063oa

Bryant, K. J. (1982). Personality correlates of sense of direction and geographical orientation. *Journal of Personality and Social Psychology*, *43*, 1318–1324.

Burgess, N. (2006). Spatial memory: how egocentric and allocentric combine. *Trends in Cognitive Sciences*, *10*(12), 551–557. doi:10.1016/j.tics.2006.10.005

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 & Virtual Environments*, *7*(2), 168–178. doi:10.1162/105474698565659

Coates, G. (1992). Program from Invisible Site-a virtual sho, a multimedia performance work presented by George Coates Performance Works. *San Francisco, CA*.

Darwin, C. (1873). Origin of Certain Instincts. *Nature*, *7*, 417–418.

Eichenberg, C., & Wolters, C. (2012). Virtual Realities in the treatment of mental disorders: A Review of the current state of research. *Virtual Reality. Rijeka, Croatia: InTech*, 35–64.

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*(6954), 184–8. doi:10.1038/nature01964

Emmelkamp, P. M. ., Krijn, M., Hulsbosch, A. ., de Vries, S., Schuemie, M. ., & van der Mast, C. A. P. . (2002). Virtual reality treatment versus exposure in vivo: a comparative evaluation in acrophobia. *Behaviour Research and Therapy*, *40*(5), 509–516.

Esch, H., & Burns, J. (1996). Distance estimation by foraging honeybees. *Journal of Experimental Biology*, *199*(1), 155–162.

Etienne, A., Maurer, R., & Saucy, F. (1988). Limitations in the assessment of path dependent information. *Behaviour*, 81–111.

Fang, L., & Loughin, T. M. (2013). Analyzing Binomial Data in a Split-Plot Design: Classical Approach or Modern Techniques? *Communications in Statistics - Simulation and Computation*, *42*, 727–740. doi:10.1080/03610918.2011.650264

Farrell, M. J., & Robertson, I. H. (1998). Mental rotation and automatic updating of body-centered spatial relationships. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(1), 227–233. doi:10.1037/0278-7393.24.1.227

Frith, H., & Gleeson, K. (2004). Clothing and Embodiment: Men Managing Body Image and Appearance. *Psychology of Men & Masculinity*, *5*(1), 40–48. doi:10.1037/1524-9220.5.1.40

Frost, P., & Warren, P. (2000). Virtual reality used in a collaborative architectural design process. In *2000 IEEE Conference on Information Visualization. An International Conference on Computer Visualization and Graphics* (pp. 568–573). IEEE Comput. Soc. doi:10.1109/IV.2000.859814

Gardony, A. L., Taylor, H. a, & Brunyé, T. T. (2014). What does physical rotation reveal about mental rotation? *Psychological Science*, *25*, 605–12. doi:10.1177/0956797613503174

Gibson, J. J. (1979). *The ecological approach to visual perception.* Psychology Press.

Goeke, C. M., König, P., & Gramann, K. (2013). Different strategies for spatial updating in yaw and pitch path integration. *Frontiers in Behavioral Neuroscience*, *7*(February), 5. doi:10.3389/fnbeh.2013.00005

Gramann, K. (2013). Embodiment of an Individual Proclivities for Egocentric and Allocentric Reference Frames. *Spatial Cognition and Computation*, *13*, 1–25. doi:10.1080/13875868.2011.589038

Gramann, K., El Sharkawy, J., & Deubel, H. (2009). Eye-movements during navigation in a virtual tunnel. *The International Journal of Neuroscience*, *119*, 1755–1778. doi:10.1080/00207450903170361

Gramann, K., Müller, H. J., Eick, E.-M., & Schönebeck, B. (2005). Evidence of separable spatial representations in a virtual navigation task. *Journal of Experimental*

Psychology. Human Perception and Performance, *31*(6), 1199–223. doi:10.1037/0096-1523.31.6.1199

Gramann, K., Müller, H. J., Schönebeck, B., & Debus, G. (2006). The neural basis of ego- and allocentric reference frames in spatial navigation: evidence from spatio-temporal coupled current density reconstruction. *Brain Research*, *1118*(1), 116–29. doi:10.1016/j.brainres.2006.08.005

Gramann, K., Onton, J., Riccobon, D., Müller, H. J., Bardins, S., & Makeig, S. (2010). Human brain dynamics accompanying use of egocentric and allocentric reference frames during navigation. *Journal of Cognitive Neuroscience*, *22*, 2836–2849. doi:10.1162/jocn.2009.21369

Gramann, K., Wing, S., Jung, T.-P., Viirre, E., & Riecke, B. E. (2012). Switching Spatial Reference Frames for Yaw and Pitch Navigation. *Spatial Cognition & Computation*, *12*, 159–194. doi:10.1080/13875868.2011.645176

Grant, S. C., & Magee, L. E. (1998). *Contributions of proprioception to navigation in virtual environments. Human factors* (Vol. 40, pp. 489–497). doi:10.1518/001872098779591296

Greenbaum, P. (1992). The lawnmower man. *Film and Video*, *9*(3), 58–62.

Haber, L., Haber, R. N., Penningroth, S., Novak, K., & Radgowski, H. (1993). Comparison of nine methods of indicating the direction to objects: data from blind adults. *Perception*, *22*, 35–47. doi:10.1068/p220035

Hegarty, M., Montello, D. R., Richardson, A. E., Ishikawa, T., & Lovelace, K. (2006). Spatial abilities at different scales: Individual differences in aptitude-test performance and spatial-layout learning. *Intelligence*, *34*, 151–176. doi:10.1016/j.intell.2005.09.005

Howarth, P. A., Sharples, S., Cobb, S., Moody, A., & Wilson, J. R. (2008). Virtual reality induced symptoms and effects (VRISE): Comparison of head mounted display (HMD), desktop and projection display systems. *Displays*, *29*(2), 58–69.

Ishikawa, T., & Montello, D. R. (2006). Spatial knowledge acquisition from direct experience in the environment. *Cognitive Psychology*, *52*(2), 93–129. doi:10.1016/j.cogpsych.2005.08.003

Johnson, R. B., & Onwuegbuzie, A. J. (2004). Mixed methods research: A research paradigm whose time has come. *Educational Researcher*, *33*(7), 14–26.

Klatzky, R. L. (1998). Allocentric and egocentric spatial representations: Definitions, distinctions, and interconnections. In *Spatial cognition* (pp. 1–17). Springer.

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. doi:10.1111/1467-9280.00058

Lawson, B. D., & Riecke, B. E. (2014). The Perception of Body Motion. In K. S. Hale & K. M. Stanney (Eds.), *Handbook of Virtual Environments: Design, Implementation, and Applications, Second Edition* (pp. 163–195). CRC Press.

Loomis, J. M., Blascovich, J. J., & Beall, A. C. (1999). Immersive virtual environment technology as a basic research tool in psychology. *Behavior Research Methods*, *31*(4), 557–564.

Malinowski, J. C. (2001). Mental rotation and real-world wayfinding. *Perceptual and Motor Skills*, *92*(1), 19–30. doi:10.2466/pms.2001.92.1.19

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

McNamara, T. P., Sluzenski, J., Rump, B., & Byrne, J. H. (2008). Human Spatial Memory and Navigation (pp. 157–178). Oxford: Academic Press.

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the “cognitive map”. *Nature Reviews. Neuroscience*, *7*(8), 663–78. doi:10.1038/nrn1932

Minsky, M. (1980). Telepresence. *Omni Magazine*, 45–52.

Mittelstaedt. (1982). Homing by Path Integration. *Proceedings in Life Sciences*.

Mittelstaedt. (1985). Analytical cybernetics of spider navigation. *Neurobiology of Arachnids*, 298–316.

Mittelstaedt, & Glasauer, S. (1991). Idiothetic navigation in gerbils and humans. *Zool. Jb. Physiol*, *95*(427-435).

Mühlberger, a, Herrmann, M. J., Wiedemann, G. C., Ellgring, H., & Pauli, P. (2001). Repeated exposure of flight phobics to flights in virtual reality. *Behaviour Research and Therapy*, *39*(9), 1033–50.

Müller, M., & Wehner, R. (1994). The hidden spiral: systematic search and path integration in desert ants, Cataglyphis fortis. *Journal of Comparative Physiology A*, *175*(5), 525–530. doi:10.1007/BF00199474

Muller, R. (1996). A Quarter of a Century of Place Cells. *Neuron*, *17*(5), 813–822. doi:10.1016/S0896-6273(00)80214-7

Ota, D., Loftin, B., Saito, T., Lea, R., & Keller, J. (1995). Virtual reality in surgical education. *Computers in Biology and Medicine*, *25*(2), 127–137.

Parkin, S. (2013). *Oculus Rift Brings Virtual Reality to Verge of the Mainstream | MIT Technology Review*. *Computing News*.

Pausch, R., Snoddy, J., & Taylor, R. (1996). Disney's Aladdin: first steps toward storytelling in virtual reality. *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, 193–205.

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

Peters, M., Laeng, B., Latham, K., & Jackson, M. (1995). A redrawn Vandenberg and Kuse mental rotations test: different versions and factors that affect performance. *Brain and Cognition*, *28*, 39–58. doi:10.1006/brcg.1995.1032

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

Psotka, J. (1995). Immersive training systems: Virtual reality and education and training. *Instructional Science*, *23*(5), 405–431. doi:10.1007/BF00896880

Riecke, B. E. (2003). *How far can we get with just visual information? Path integration and spatial updating studies in Virtual Reality*. Universit\"atsbibliothek.

Riecke, B. E. (2008). Consistent Left-Right Reversals for Visual Path Integration in Virtual Reality: More than a Failure to Update One's Heading? *Presence: Teleoperators & Virtual Environments*, *17*(2), 143–175. doi:10.1162/pres.17.2.143

Riecke, B. E. (2012). Are Left-right Hemisphere Errors in Point-to-origin Tasks in VR Caused by Failure to Incorporate Heading Changes? In *Proceedings of the 2012 International Conference on Spatial Cognition VIII* (pp. 143–162). Berlin, Heidelberg: Springer-Verlag. doi:10.1007/978-3-642-32732-2_9

Riecke, B. E., & Bülthoff, H. H. (2004). Spatial updating in real and virtual environments: contribution and interaction of visual and vestibular cues. In *Proceedings of the 1st Symposium on Applied perception in graphics and visualization* (Vol. 1, pp. 9–17). ACM.

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. doi:10.1007/s00426-006-0085-z

Riecke, B. E., Heyde, M. V. ., & Bülthoff, H. H. (2001). How real is virtual reality really? Comparing spatial updating using pointing tasks in real and virtual environments. *Journal of Vision*, *1*(3), 321–321. doi:10.1167/1.3.321

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

Riecke, B. E., Schulte-Pelkum, J., & Bülthoff, H. H. (2005). Perceiving Simulated Ego-Motions in Virtual Reality - Comparing Large Screen Displays with HMDs. In *Proceedings of the International Society for Optical Engineering* (Vol. 5666, pp. 344–355). Citeseer. doi:10.1117/12.610846

Riecke, B. E., Sigurdarson, S., & Milne, A. P. (2012). Moving through virtual reality without moving? *Cognitive Processing*, *13*(1), 293–297.

Riecke, B. E., Veen, H. A. H. C., & Bülthoff, H. H. (2002). Visual homing is possible without landmarks: A path integration study in virtual reality. *Presence: Teleoperators & Virtual Environments*, *11*(5), 443–473.

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. *2007 IEEE Virtual Reality Conference*, 3–10. doi:10.1109/VR.2007.352457

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. doi:10.1037/0278-7393.15.6.1157

Rieser, J. J., Guth, D. A., & Hill, E. W. (1982). Mental processes mediating independent travel: Implications for orientation and mobility. *Journal of Visual Impairment & Blindness*, *76*, 213–218.

Rothbaum, B. O., Hodges, L. F., Kooper, R., Opdyke, D., Williford, J. S., & North, M. (1995). Effectiveness of computer-generated (virtual reality) graded exposure in the treatment of acrophobia. *The American Journal of Psychiatry*, *152*(4), 626–8.

Ruddle, R., & Lessels, S. (2006). For efficient navigational search, humans require full physical movement, but not a rich visual scene. *Psychological Science*, *17*(6), 460–5. doi:10.1111/j.1467-9280.2006.01728.x

Scozzari, S., & Gamberini, L. (2011). Virtual Reality as a Tool for Cognitive Behavioral Therapy: A Review. *Advanced Computational Intelligence Paradigms in Healthcare 6. Virtual Reality in Psychotherapy, Rehabilitation, and Assessment*, 63–108.

Sigurdarson, S., Milne, A. P., Feuereissen, D., & Riecke, B. E. (2012). Can physical motions prevent disorientation in naturalistic VR? In *Virtual Reality Short Papers and Posters (VRW), 2012 IEEE* (pp. 31–34).

Silverman, I., & Eals, M. (1992). Sex differences in spatial abilities: Evolutionary theory and data. In *Portions of this paper were presented at the meetings of the International Society for Human Ethology in Binghamton, NY, Jun 1990, the Human*

Behavior and Evolution Society in Los Angeles, CA, Aug 1990, and the European Sociobiological Society in Prague, C (Vol. The adapte, pp. 533–549).

Skolnick, N. J., Ackerman, S. H., Hofer, M. A., & Weiner, H. (1980). Saccades are spatially, not retinocentrically, coded. *Science*, *208*, 6.

Smith, S., & Ericson, E. (2009). Using immersive game-based virtual reality to teach fire-safety skills to children. *Virtual Reality*, *13*(2), 87–99. doi:10.1007/s10055-009-0113-6

Sproll, D. (2013). *Influence of Ethnicity, Gender and Answering Mode on Reference Frame Selection for Virtual Point-to-Origin Tasks*. Bachelor's thesis. Universität Osnabrück & Simon Fraser University.

Stanney, K. M., Mourant, R. R., & Kennedy, R. S. (1998). Human Factors Issues in Virtual Environments: A Review of the Literature. *Presence: Teleoperators and Virtual Environments*, *7*, 327–351. doi:10.1162/105474698565767

Steuer, J. (1992). Defining Virtual Reality: Dimensions Determining Telepresence. *Journal of Communication*, *42*(4), 73–93. doi:10.1111/j.1460-2466.1992.tb00812.x

Sutherland, I. (1965). The ultimate display. *Multimedia: From Wagner to Virtual Reality*.

Thinus-Blanc, C., & Gaunet, F. (1997). Representation of space in blind persons: vision as a spatial sense? *Psychological Bulletin*, *121*(1), 20–42.

Vandenberg, S. G., & Kuse, A. R. (1978). Mental rotations, a group test of three-dimensional spatial visualization. *Perceptual and Motor Skills*, *47*, 599–604. doi:10.2466/pms.1978.47.2.599

Virvou, M., Manos, C., Katsionis, G., & Tourtoglou, K. (2002). VR-ENGAGE: A virtual reality educational game that incorporates intelligence. In *Proceedings of IEEE international conference on advanced learning technologies* (pp. 16–19).

Wan, X. I., Wang, R. F., & Crowell, J. A. (2009). Spatial updating in superimposed real and virtual environments. *Attention Perception Psychophysics*, *71*, 42–51.

Wang, R. F. (2004). Between reality and imagination: when is spatial updating automatic? *Perception & Psychophysics*, *66*(1), 68–76.

Wang, R. F., & Brockmole, J. R. (2003). *Simultaneous spatial updating in nested environments. Psychonomic bulletin review* (Vol. 10, pp. 981–986).

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

Wang, R. F., & Spelke, E. S. (2002). Human spatial representation: Insights from animals. *Trends in Cognitive Sciences*, *6*(9), 376–382.

Wartenberg, F., May, M., & Péruch, P. (1998). Spatial orientation in virtual environments: Background considerations and experiments. In *Spatial cognition* (pp. 469–489). Springer.

Wehner, R., & Müller, M. (2010). Piloting in desert ants: pinpointing the goal by discrete landmarks. *The Journal of Experimental Biology*, *213*(24), 4174–4179.

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

Wolbers, T., & Hegarty, M. (2010). What determines our navigational abilities? *Trends in Cognitive Sciences*. doi:10.1016/j.tics.2010.01.001

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.