The Effect of Transition Type in Multi-View 360° Media

Andrew MacQuarrie and Anthony Steed

Abstract—360° images and video have become extremely popular formats for immersive displays, due in large part to the technical ease of content production. While many experiences use a single camera viewpoint, an increasing number of experiences use multiple camera locations. In such multi-view 360° media (MV360M) systems, a visual effect is required when the user transitions from one camera location to another. This effect can take several forms, such as a cut or an image-based warp, and the choice of effect may impact many aspects of the experience, including issues related to enjoyment and scene understanding. To investigate the effect of transition types on immersive MV360M experiences, a repeated-measures experiment was conducted with 31 participants. Wearing a head-mounted display, participants explored four static scenes, for which multiple 360° images and a reconstructed 3D model were available. Three transition types were examined: teleport, a linear move through a 3D model of the scene, and an image-based transition using a Möbius transformation. The metrics investigated included spatial awareness, users' movement profiles, transition preference and the subjective feeling of moving through the space. Results indicate that there was no significant difference between transition types in terms of spatial awareness, while significant differences were found for users' movement profiles, with participants taking 1.6 seconds longer to select their next location following a teleport transition. The model and Möbius transitions were significantly better in terms of creating the feeling of moving through the space. Preference was also significantly different, with model and teleport transitions being preferred over Möbius transitions. Our results indicate that trade-offs between transitions will require content creators to think carefully about what aspects they consider to be most important when producing MV360M experiences.

Index Terms—H.5.1 [Information interfaces and presentation]: Multimedia Information Systems – Artificial, augmented, and virtual realities

1 Introduction

Due in part to the rising quality of 360° cameras, captured 360° media is an increasingly appealing way to create immersive experiences. Captured 360° media, however, is generally fixed viewpoint i.e. only the three degrees of freedom associated with orientation are available to the user to explore the scene. While free-viewpoint 360° media is being investigated (e.g. [9, 11]), there are still many technical challenges to be overcome. Multi-view 360° media (MV360M) – in which 360° views are captured from multiple locations – may offer a partial solution to the issues associated with fixed-viewpoint 360° media. This is achieved by allowing users to view the space from multiple perspectives. Systems have existed for some time that allow the exploration of spaces by transitioning between 360° images [1, 10], and video-based MV360M experiences are becoming more common [3, 29].

In MV360M systems, a visual effect is required to transition the user from one camera location to another. This effect may have a significant impact on user experience. While multiple transitions are available from standard film production, such as wipe, dissolve, fade, etc., there are aspects inherent in immersive MV360M that require special consideration. A fade to black in an immersive display, for example, is the equivalent of the world suddenly going dark, and with no visual features available, this transition might result in discomfort or disorientation.

Likewise, while work has been done to explore the impact of transition types in immersive experiences – particularly in the virtual reality (VR) locomotion literature – certain attributes of MV360M make such systems inherently different from most real-time rendered experiences. For example, as each view in MV360M requires a physical camera to have been placed at that location, such locations tend to be limited in number. The limited number and fixed-position nature of the views available may have a detrimental effect on a user's ability to understand the scene. This may mean that additional benefit could be gained from transition types that provide further information about the spatial layout of the environment over the same transition in real-time rendered experiences.

In our work, we explored the impact of the transition type when viewing image-based MV360M. A repeated-measures experiment was conducted with 31 participants. Wearing a head-mounted display (HMD), participants navigated through four static scenes, initiating transitions using a tracked hand controller. Three transition types were examined: *teleport*, *model* and *Möbius*. The teleport transition moves the user instantaneously to their selected location while leaving their orientation unaltered. The model transition moves the user linearly through a reconstructed 3D model of the scene. While parallax cues inherent in the model transition are likely to provide users with the most complete impression of the scene, such models are expensive and labour-intensive to produce. We propose the Möbius transition as a possible middle ground between the teleport and model transitions. The Möbius transition is an image-based transformation that gives the impression of movement between panoramas using a zoom effect. The motion cues provided by this effect may help to improve understanding of the scene and the transition, while as an image-based transition there is no requirement for a 3D reconstruction of the scene to be available.

The metrics investigated were spatial awareness, users' movement profiles, transition preference and the subjective feelings of moving through the space, disorientation, dizziness, and naturalness. Our results indicate that trade-offs between transitions will require content creators to think carefully about what aspects they consider to be most important when producing MV360M experiences. Additionally, unexpectedly poor spatial awareness results across all conditions may indicate that MV360M experiences do not facilitate scene understanding.

2 Related Work

2.1 360° Media

Most 360° media is recorded by capturing several overlapping camera views of the same scene. These views are then stitched together in software. This creates a "viewing sphere" – a 360° view as seen from a single location. While these 360° views are often captured using rigs of several cameras, we will use the generic term "camera" to refer to a single or multiple capture cameras at a single location.

MV360M can help to facilitate the exploration of a captured scene

- *Andrew MacQuarrie is with University College London. E-mail: andrew.macquarrie.13@ucl.ac.uk.*
- *Anthony Steed is with University College London. E-mail: a.steed@ucl.ac.uk.*

Manuscript received 11 Sept. 2017; accepted 8 Jan. 2018.
Date of publication 19 Jan. 2018; date of current version 18 Mar. 2018.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TVCG.2018.2793561

by allowing users to view the space from multiple locations. Such systems can use 360° images or video. Various VR experiences have made use of these content types in different ways. QuickTime VR allowed the user to explore a scene by navigating between 360° images captured from different locations [10]. This type of content has become commonplace, for example in systems such as Google Street View [1]. Video-based MV360M has historically been less popular, for example due to typical bandwidth limitations; it is, however, becoming more common. Video-based MV360M has been live streamed from events including award shows [3] and concerts [29], allowing users to choose between views in real-time.

2.2 Locomotion and Spatial Awareness

Similarly to VR locomotion, MV360M allows users to explore virtual spaces. This is a useful parallel to draw, as the VR locomotion literature has established several metrics to evaluate different techniques. Previously investigated locomotion techniques include teleportation [7], auto-locomotion [34], walking-in-place [36], redirected walking [30] and the use of a visual metaphor such as a portal [17].

While transitions through MV360M can be framed as a VR locomotion task, there are certain differences that are important. Often, locomotion techniques are assessed on metrics such as accuracy of positioning, speed, number of collisions, and ease-of-control (e.g. as in [36]). Such metrics are not necessarily appropriate for examining different transitions. For example, the accuracy of positioning cannot reasonably be examined, as the only available locations are predetermined by the camera positions, and it is usually impossible to miss them due to the nature of the interface.

There are several metrics in the VR locomotion literature, however, that are useful for exploring the effects of transition types. Aspects such as the transition's effect on the user's spatial awareness is of importance, as content producers may wish to understand how their choice of transition will affect a user's understanding of the captured space. This may be of particular importance in MV360M, as the lack of parallax cues from head movement may have a detrimental effect on a user's spatial awareness over the six degrees of freedom generally associated with HMD experiences. Bowman et al. looked at the effect of transition types on spatial awareness in their work on viewpoint control techniques [6]. They concluded that teleportation transitions produced poorer spatial awareness than moving through the space, as assessed by the time taken to visually find a previously seen object.

Metrics from the spatial awareness literature have also been employed. Pointing tasks, in which the user is asked to indicate the direction of a previously seen object that is no longer visible, have been used in the spatial awareness literature to gauge participants' understanding of both physical buildings [35] and large outdoor spaces [28]. Pointing tasks similar to these have been used in the VR locomotion literature. In work by Bowman et al., a pointing task was used to evaluate a user's ability to maintain spatial orientation while navigating through virtual corridors [5]. Their results indicated that locomotion techniques in which the user did not physically move their body still allowed them to maintain spatial orientation. Recently, Sargunam et al. used a pointing task to evaluate the effect of amplified and guided head rotations on spatial awareness [33]. Their results indicated that guided head rotations may negatively impact spatial awareness, but only found the effect to be significant for participants with significant gaming experience. As well as objective measures, subjective aspects such as the naturalness of the transition (e.g. [36]) and user preference (e.g. [7]) will have an impact on the user's experience.

As discussed by Bowman et al. in their work on VR locomotion, there is an important distinction between locomotion and navigation [6]. Navigation is a complex area that incorporates many cognitive processes. While navigation is undeniably an important concept when exploring a space in VR, like Bowman et al. we do not attempt to address the underlying processes involved, although work has previously been done in this area [12].

2.3 Transitions

The creation of transitions between panoramic images has been studied for some time. McMillan and Bishop first proposed techniques for creating novel views by interpolating between panoramic images captured from cameras with a small baseline [25]. Morvan and O'Sullivan continued this work to extend the required baseline [26]. In their work, parallax was faked by billboarding foreground objects using occluder masks, using a technique similar to Tour into the Picture [19]. Morvan and O'Sullivan also used laser scanners to create accurate models of scenes; however producing such models required expensive specialist hardware and was labour intensive, so this technique was not used in their final experiments. Morvan and O'Sullivan also conducted a user study to establish transition preference between faked parallax, a dip to black (fade) and a cross dissolve (blend). They concluded that blending was always preferred over fading, and that fake parallax was generally preferred over blending.

These works, however, relate to panoramic media exploration in a desktop setting. Viewing panoramic media in an immersive display such as a HMD is substantially different. For example, the user study by Morvan and O'Sullivan did not explore "cutting" (an instantaneous transition) as it was not considered to be "well suited to continuous navigation". In an immersive context, instantaneous transitions are frequently used for navigation, largely due to the effect of vection on simulator sickness [18]. Additionally, in the work by Morvan and O'Sullivan, the orientation of the view during transitions was predetermined, which is not conducive to a HMD experience where the view is usually determined by the orientation of the HMD.

2.4 Simulator Sickness

Simulator sickness is of particular importance when dealing with transitions in a HMD. This is due to the fact that simulator sickness can be induced through vection in a VR display [18], and vection is a necessary component of some transition types. There have been studies that indicate that a user is less likely to experience simulator sickness if they can control or anticipate the motion [34]. Additionally, there is evidence that most users become less susceptible to particular movements with repeated exposure [20]. A common way to measure simulator sickness in VR experiments is via self assessment, using the Simulator Sickness Questionnaire (SSQ) [21].

3 Experimental Design

We conducted a repeated-measures user study to evaluate the effect of transition type when exploring a scene captured in 360° from multiple locations. At each location, participants could look around naturally via a tracked HMD, using the three degrees of freedom associated with orientation. Participants could move around inside the scenes by selecting different camera locations using a position-tracked, hand-held input device. While there were multiple buttons and triggers on the input device, they all performed the same action and users were free to use whichever button felt most comfortable. When a user chose to move to another location, they were transitioned from their current location to their selected location via one of three transition types, each of which is described in section 3.2. Generally, the transition type was selected randomly by the system. The only time when the transition type was not randomly selected was before the pointing task or preference questions; at these times the transition type was counterbalanced. These tasks and questions are described fully in section 3.3, while our counterbalancing strategy is described in section 3.6.

3.1 Stimuli

As will be discussed in section 3.2, the model transition required a 3D reconstruction of the scene. As a result of this, the available stimuli was limited to static scenes. The stimuli was provided by Matterport, a VR capture company whose cameras record RGB-D data in 360°. Using this data, Matterport construct a textured 3D model of the scene. This provides 4k 360° images from set locations, as well as a 3D model of the scene. The quality of these 3D models is consistent with structured light scanning, i.e. they contain some holes and do not accurately

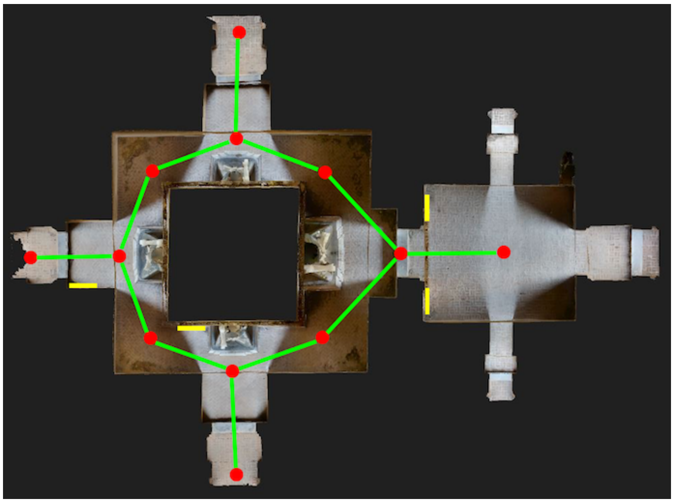

Fig. 1: A map of the temple stimulus. Camera locations are shown as red nodes, while green edges indicate available transitions between locations. Yellow lines show the locations of targets in this scene.

Table 1: Stimuli statistics. Stimuli are listed in the order they were presented to participants.

Stimulus	Camera locations	Transition edges	Targets
Gallery	4	3	2
House upstairs	11	10	6
Temple	12	12	4
House downstairs	11	12	6

capture fine details. The 3D model for one of our stimuli is available to view online [16].

In total, three models were used: a gallery, a Buddhist temple, and a house. As the house was set over two floors, this model was split into two independent scenes. This provides a total of four stimuli: *gallery*, *temple*, *house upstairs* and *house downstairs*. A map of the temple scene is shown in Figure 1, while statistics for all stimuli are shown in Table 1.

The stimuli was limited to indoor scenes. In general, scenes were highly occluded, meaning only a small number of cameras were in the line-of-sight of any other camera. As our setup only allows transitions to cameras that are in the line-of-sight of the current camera location, this required participants to navigate around the scene, transitioning from camera to camera, in order to complete the required tasks. This procedure is described fully in section 3.6. Camera locations were limited to the minimum number possible, such that the scene was covered by a connected graph of transitions.

3.2 Transition Types

Three transition types were used: teleport, model and Möbius. The teleport transition was instantaneous, while both the model and Möbius transitions took six seconds to complete. Each of these transitions is now described in detail.

Teleport: When a participant chose to move to a new location, and the transition was determined by the system to be a teleport, the participant was moved instantaneously from their original location to their selected location. On arrival at this new location, their orientation in the virtual space was the same as it had been at their start location. To achieve this, the image the user was seeing was instantaneously swapped from the panorama captured at their original location to the panorama captured at their selected location. Then, as in the other transitions, scene elements such as the available locations to move to next were updated to be consistent with the new panorama.

Model: In the model transition, the user moves through a 3D model of the scene. This transition, therefore, requires a reconstructed 3D

Fig. 2: The 3D model transition. First, a blend is performed between the panoramic image for the original location (frame a) and the 3D model (frame b). The user is then moved linearly through the scene (frame c). Finally, a blend between the 3D model and the panoramic image for the new location is performed (frame d).

model of the scene. While movement easing types were explored, simulator sickness appears least severe when there is minimal changes in velocity [4]. Following advice from Oculus that it is the duration of velocity change that should be minimised, a linear movement with infinite acceleration and deceleration was used [27]. As the user sees a panoramic image when not moving between locations, a linear interpolation was used to blend between the panoramic images and the 3D reconstruction of the scene. The blend was necessary as the 3D model was not completely consistent with the panoramic images, for example the lighting was often noticeably different. This would likely always be the case for MV360M. Even if all cameras could have matching settings such as exposure, the location of the camera impacts aspects such as specular highlights, so the lighting would not be consistent between 360° cameras capturing the same scene. Blending between the panorama of the current location and the model lasted 0.5s, the linear movement from the current location through the 3D model to the new location lasted 5s, and the blend from the model to the panorama of the new location lasted 0.5s, resulting in a total transition time of 6s. Frames from this transition type can be seen in Figure 2.

Möbius: The Möbius transition is an image-based transition, in which a transformation is applied to give the impression that the user is moving from the panorama of the current location to the panorama of the next locations. This is achieved by "zooming in" to both panoramic images. While zooming in is a common technique in standard format media production that gives the illusion of getting closer to something, its application in panoramic media is complicated by the spherical nature of the imagery. In order to zoom in to a specific point in a spherical image, the rest of the content cannot be cropped, but must instead be compacted towards the zoom's antipodal point.

The Möbius transformation was proposed by eleVR and Henry Segerman as a technique to allow zooming for panoramic media [14]. The Möbius transformation is a conformal mapping, in that it preserves local angles (for a review, see [2]). It can be used to enlarge the image towards the zoom point, while reducing the size of elements towards the zoom's antipodal point. This technique can be used to give the impression of moving towards the zoom point, although deformations to the space mean the effect does not appear natural.

There are many possible ways to incorporate the Möbius transformation into a transition. In our implementation, the user selects a new location to move to. We call the panoramic image for the current location L_c, and the panoramic image for the selected next location as L_n. If a ray was cast from the current location to the new location in 3D space, the point of intersection on the viewing sphere of L_c becomes the "zooming point" for the Möbius transformation, referred to as P_{zc}

here. Likewise, the zooming point for L_n is where the same ray would intersect L_n, referred to as P_{zn}.

First, a "zoomed out" version of L_n expands as a circle at P_{zc} until it reaches approximately 38% of the height of L_c when viewed in equirectangular form. A value of 38% was chosen for aesthetic reasons. The Möbius transformation is then applied to "zoom into" L_c and the zoomed out version of L_n. Assuming that the viewer is facing towards the new location, this means L_c collapses behind the viewer, while L_n expands over them. In our implementation, all transitions were generated as video files in a preprocessing step. This video file was then played back when a transition was initiated, using the PopMovie video plugin for Unity. Due to larger videos causing unacceptable levels of lag in the rendering, the videos were downscaled and played at a resolution of 2048x1024. During a transition, the video was started and blended in over 0.5s, continued playing for 5s, and then blended out over 0.5s, resulting in a total transition time of 6s. This effect can be seen in Figure 3.

The effect produced by the Möbius transformation is difficult to describe accurately, and we would encourage readers to watch our videos of the transition in practice. We have recorded a complete user journey through the gallery stimulus [23], and made available a 360° video showing two Möbius transitions in the temple stimulus [22]. These videos are also included in the supplementary materials. The code used to generate our Möbius transitions is available online [24].

3.3 Hypotheses

3.3.1 Spatial Awareness

As in the related work discussed in section 2.2, spatial awareness was measured using a pointing task. Participants were asked to point at a known location in the scene that was no longer visible using a tracked hand controller. We refer to this task as the "pointing task". Both the error angle and the time to complete the task were examined. The error angle was defined as the angle between the user's pointing ray and the ray from the centre of the hand controller to the centre of the target in question. This angle was calculated on a 2D plane as seen from above, i.e. the elevation components were discarded for both rays.

It was expected that the teleportation transition would provide the poorest spatial awareness, while the model transition would provide the best. As the Möbius transition provides some movement cues, it was expected that this transition would produce a spatial awareness result somewhere between the model and the teleportation transitions.

H1: It was hypothesised that the transition type would have an effect on spatial awareness.

3.3.2 Subjective Measures

H2: It was hypothesised that the transition type would have an effect on participants' subjective experience of moving through the space, dizziness, disorientation and naturalness.

These metrics were assessed by asking the participant to verbally provide a rating from one to five, where one meant "not at all" and five meant "extremely". These four questions were, for the transition that they just saw: how much did they feel that they were moving through the space (moving); to what extent did they feel disoriented (disoriented); how dizzy did they feel (dizzy); how natural did the transition feel (naturalness). Naturalness in this context was described to participants as, "how organic and close to real life" the transition felt. Participants were asked to consider each transition on its own. For example, when asking about "disorientation", participants were told that this was not about their general sense of confusion about the scene, but whether or not they had been disoriented by that transition.

Moving: an important characteristic of a transition is whether or not a user feels as if they are moving through the scene. As the teleportation transition is instantaneous, it was expected that users would not feel that they are "moving through the space" during this transition. It was expected that users would strongly feel that they are moving through the space during the 3D model transition. As the Möbius transition provides some movement cues, it was expected that participants would feel somewhat as if they are moving through the space during this transition.

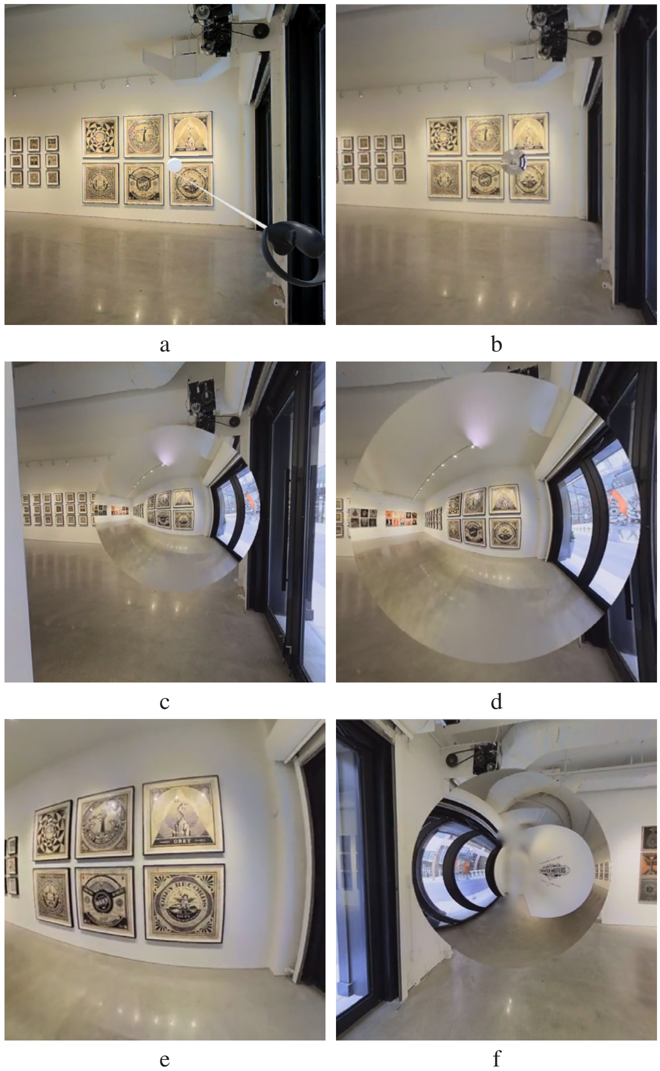

Fig. 3: Frames from the Möbius transition. Frame a: the user initiates a transition by selecting the new location's marker using the input device. Frames b-c: a zoomed out version of the new location's panorama expands into view. Frames d-e: the Möbius transformation is applied to both panoramas, creating a zooming effect. Frame f: looking backwards, the previous location collapses behind the user.

Dizzy: in direct contrast to a participant feeling like they are moving through the space, it was expected that teleportation would produce a low rating for dizziness, while Möbius and 3D model transitions would produce higher ratings.

Disoriented: as the teleportation transition has previously been shown to disorient users [6], it was expected that this transition would produce a higher subjective rating for disorientation. The model transition was expected to produce the lowest disorientation result, while the effect of the Möbius transition was unclear.

Natural: it was expected that the Möbius and teleportation transitions would be rated poorly for naturalness, while the 3D model transition would receive a higher rating.

Additionally, participants' preferences for transition types were explored. It was unclear what transitions users would like most. This was measured by asking participants to state a binary preference between the last two transitions they saw. A preference value was taken for each possible pairing and ordering of transition types.

H3: It was hypothesised that there would be a difference in participants' preferences for transition types.

3.3.3 Movement Profile

H4: It was hypothesised that the transitions type would have an effect on the movement profile of a user.

Movement profile was measured by examining the time taken to initiate the next transition, following the completion of the previous transition. As teleportation is more likely to disorient participants, it was expected that it would take longer for a user to initiate the next transition following the teleport transition than following the model or Möbius transitions.

3.4 Experimental Setup

Participants wore an Oculus CV1 HMD. The CV1 was driven by a Windows 10 desktop PC with an Intel i7-6700 CPU running at 3.4GHz with 32GB of RAM. The video card in use was a NVIDIA GeForce GTX 1080. The software was implemented using the Unity game engine, version 5.6.2f1.

As the captured 360° images must be viewed from the centre of the viewing sphere, only the three degrees of freedom associated with orientation were available to users through their head-tracked movements. This meant there was no visual feedback available to participants in regards to their physical position. As a result of this – coupled with the vection caused by some of the transitions – it was decided that having a participant stand during the study could be unsafe. To ensure safety, participants were seated while wearing the HMD. A swivel chair was used to allow the participant to rotate freely, while the HMD was suspended from the ceiling to avoid any movement limitations that would otherwise have been caused by the cable.

As the experiment required the use of a handheld, tracked input device, an Oculus Touch controller was used. The right-hand controller was used, however as our interface did not require use of the trigger buttons, participants could hold the controller in their preferred hand. Use of the Oculus Touch controller required 360° positional tracking. To facilitate this, three Oculus Sensors were placed in a triangle around the swivel chair facing inwards. This provided accurate positional tracking for the hand controller, as well as drift correction for the rotation tracking of the HMD. When wearing the HMD, participants could see a virtual representation of the hand controller. This virtual representation included a ray, to make it easy for participants to identify the exact direction the hand controller was pointing. The virtual hand controller maintained the same relative position from the user's head as in the real world, even though the HMD position was not used to update the virtual head position inside the viewing sphere.

Participants were shown four scenes, each of which was captured from multiple locations. Other available locations were represented to the user visually as a sphere, floating at eye height at the location in 3D space that it represented. Pointing the hand controller near the sphere caused the sphere to glow, as shown in figure 3a, giving visual feedback to the user that the location was selected. When the user pressed a button on the hand controller, they were transitioned from their current location to their selected location via one of the three transition types described in section 3.2. Only other locations in the line-of-sight of the current location were available at any time.

Throughout the study, participants were asked to find targets. Targets were brightly colored squares, placed against walls or columns inside the scene at eye height. An example of such a target is shown in Figure 4a. These targets provided a catalyst for exploration, as well as easily identifiable reference points for the pointing task, as discussed later in section 3.6.

While a 3D model of the scene was available, it was not photo realistic and contained visual artefacts. The 360° images of the scene were not stereoscopic. To keep the experience consistent, all visuals were presented monoscopically. This included the spheres that represented other locations, the colored targets, and the 3D model during the model transition. The spheres that represented other locations and the colored targets were only visible when the user was static, and were disabled during transitions.

a b

Fig. 4: A green target is found by the user (frame a). Two transitions of the same type later, the world fades to the grid environment (frame b) and the user points to where they believe that most recently seen target is in relation to their current location in the scene.

3.5 Participants

The study was approved by the UCL Research Ethics Committee. All participants were recruited via a participant pool website. Thirty-three participants took part, however data from two were excluded as they did not complete the trial. For one this was due to time limitations, and the second withdrew following discomfort from simulator sickness. Of the remaining 31 participants, 18 were female and 13 male. The mean age of participants was 27.68 (SD = 7.951).

3.6 Experimental Procedure

Participants were asked to read an information sheet, as well as completing a pre-experiment questionnaire and SSQ. The experimental procedure was then explained to them. As simulator sickness was of particular concern in this experiment due to the large amount of vection involved, the risks of simulator sickness were covered in detail, as was the participant's right to stop at any time.

Participants were asked to sit on a swivel chair, and shown the HMD and the hand controller. Participants then put on the HMD, and were shown a test scene to familiarise them with the equipment and the procedure. A 360° image of a room was shown, and the participant encouraged to rotate in their chair to view the entire room. They were then asked to find a green target, and then hit this target (point at it with the hand controller and press a button). This caused the target to disappear, and another target to appear elsewhere in the room. The participant was then directed to find this new target.

Once the second target had been located and hit, there was a delay of a few seconds before the world faded to a grid environment. In this environment, the entire world is faded out. A grid is faded in at ground level to allow the participant to maintain awareness of their orientation, as shown in Figure 4b. The participant was told that when this grid environment appeared, they would be asked to point with the hand controller at where the most recently seen target was in relation to their current location in the scene. At this time, they were asked to point towards the most recently seen target and press a button. When a button was pressed, the grid environment faded out and the virtual scene was faded back in.

Participants were then introduced to the first scene, and instructed on how to transition from one location in the scene to another. They were then asked to move around the space, moving from location to location, looking for and hitting targets.

The experiment proceeded in sets of two targets, called A and B in this example. First, a participant was asked to find target A. Each target was identified by a color, and no participants had any form of colorblindness or experienced any difficulty in identifying targets. An example of such a target can be seen in Figure 4a. Once target A was found, they would be instructed to find target B. To ensure comparability between participants, target A would only be visible from a single location. Although it was not stipulated to the participant, after hitting target A, the available locations were restricted, to ensure the participant had to follow a set route. After moving two locations - in which the participant was shown the same transition type - the world faded to the grid environment, as shown in Figure 4b. The participant was then asked to perform the pointing task, i.e. point at where target A was from their current location and press a button. The error angle was then recorded. After completion of the pointing task, the world was faded back in and the user continued looking for target B. After locating target B, they were instructed to return to target A in as few transitions as possible. Returning to target A – referred to as the *returning phase* – allows for analysis of a natural user movement pattern without interference from the target search i.e. when returning to target A, the user generally knows where they are going, and are therefore not scanning the space for targets.

In total there were nine such pairs of targets across the four scenes, resulting in nine pointing tasks. In order to balance any difference in pointing task difficulty between transition type conditions, the order of transitions shown before each pointing task was counterbalanced. With three transitions under test (3 = 3D model, T = teleportation, M = Möbius), six orderings were possible (3TM, 3MT, T3M, TM3, M3T, MT3). Each participant saw a single ordering of transitions – for example a participant assigned to the first ordering would have performed their nine pointing tasks after seeing transitions in the order 3TM3TM3TM.

At times the experimenter would ask the participant a question about their subjective experience of the transitions. Five such questions were possible. Four of these questions – moving, dizziness, disorientation and naturalness, as described in section 3.3.2 – required the participant to provide a rating from one to five. Before a transition, the experimenter would get the system to select one of these four questions. In order to balance any ordering or possible scene effects, the system selected these questions randomly. Following the transition, if the question had been asked before for that transition type, the question was skipped. If the question had not been asked before, the experimenter would pause the environment (the environment faded out slightly, and actions by the participant were disabled) and orally ask the participant to rank the transition for that metric. Once the participant had provided their response orally, the experimenter would unpause the environment and the participant would continue searching for the next target. If the experimenter accidentally asked the same question twice for one transition type, the mean of those ratings was used.

The fifth possible question was for the participant to specify which of the last two transitions they preferred. In a pilot study, it became clear that participants were unable to remember the second-to-last transition with enough clarity to provide an accurate comparison. To ensure that the participant was able to provide an answer, they were alerted two transitions in advance that the experimenter was going to ask their preference, allowing them to be actively comparing them. The experimenter also pressed a button, ensuring the system would show two different transitions. The pairs of transitions shown to the user were programmed to ensure that the participant saw each possible pairing of transitions in all orders. With three transition types, this meant the user provided a preference for all six possible pairings.

The experimenter determined when to ask these questions through observation of the participant's position, and attempted to avoid asking questions at any point when the question would interfere with other metrics. For example, if the participant was one location away from a target, the experimenter would not ask the user which of the next two transitions they preferred, as at the next location the participant may have found the target and initiated a pointing task. Likewise, the experimenter avoided asking questions during the returning phase, as this may have affected the participant's movement profile.

Due to the amount of vection involved, participants took a five minute break between each scene. This was done to reduce the cumulative effects of simulator sickness. During this break, participants were asked to complete a pen-and-paper map-placement task. This task was intended to provide a general idea of a participant's spatial awareness of the scene. Participants were asked to mark camera and target locations on a map of the environment they had just seen. Following the final scene, participants were asked to complete an SSQ again. Participants were then debriefed, and given £10 compensation for taking part.

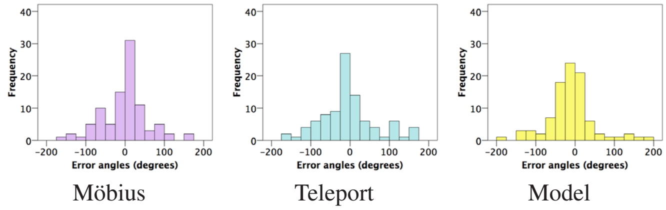

Fig. 5: Histograms of all pointing task results, including the angle's sign, for each transition type.

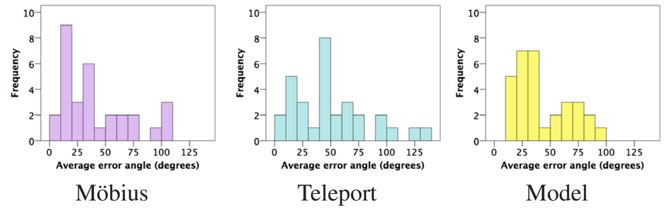

Fig. 6: Histograms of mean pointing task results for each transition type.

4 Results

4.1 Spatial Awareness

4.1.1 Error Angle

The error angle was defined as the angle between the user's pointing ray and the ray from the centre of the hand controller to the centre of the target in question, calculated on a 2D plane as seen from above. When all pointing task error angles – including their sign – are shown in a histogram, the distribution shows the expected peak near 0°, as shown in Figure 5. Although this data roughly follows a bell shape, analysis with Shapiro-Wilk indicates that the data cannot be considered normally distributed, most likely due to the frequency and spread of extreme data points. Pointing task error angles were often extremely high, indicating that participants struggled with this task. Indeed, participants frequently had error angles over 90°. For a breakdown of error angles by pointing task, please see the supplementary materials. Erratic, large error angles produced high standard deviations, making statistical analysis challenging.

When performing statistical analysis, the average of all three pointing task error angles was taken for each transition type. Only the magnitude of the error angle was used. Large error angles, however, were frequent in the data. When the average of three pointing tasks was taken, this often resulted in unrepresentative values. For example, a participant with error angles 11.7°, 13.1° and 115.7° following teleportation transitions results in an average of 46.8°. A participant having a large error angle was so frequent in the data that taking the average of three pointing tasks for each transition type resulted in the data appearing multimodal, as shown in Figure 6.

As the data is non-normal, a non-parametric test was required. A Friedman test was run to determine if there were differences in absolute pointing task error angles following three different transition types. Error angle increased from model (Mdn = 32.38°), to Möbius (Mdn = 35.18°), to teleport (Mdn = 46.81°), but the differences were not statistically significant, $\chi^2(2) = 2.774$, p = .250.

Post Hoc Analysis

It is clear that participants struggled with this task, and the data does not fit expectations. The median values for error angles appear quite different between transition types, with teleport producing a median value 45% larger than the model transition. The high variability and large number of extreme values, however, make analysis difficult. As a result, we examined the data for possible post hoc analysis techniques.

The average pointing task data contained no outliers, with an outlier being defined as a data point beyond 1.5 times the interquartile range

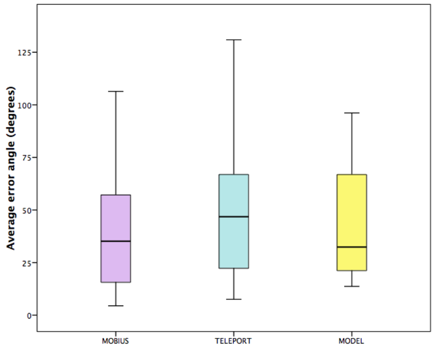

Fig. 7: Boxplot of average pointing task results for each transition type.

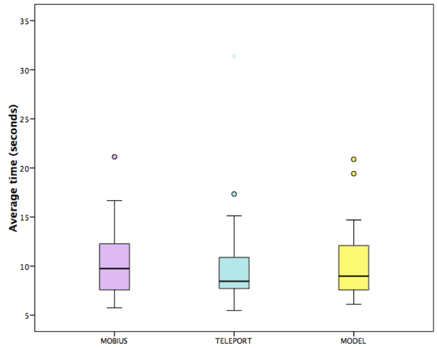

Fig. 8: Boxplot of average times between fading to grid environment and completion of pointing task.

(IQR). This can be seen in the boxplot shown in Figure 7, with whiskers representing 1.5 times the IQR. This is due to the high frequency with which participants produced high error angles, resulting in a large IQR, and means the data is not suitable for filtering out outliers.

An alternative approach is to define a sensible cutoff value, and remove all values above that cutoff across all transition types. The issue with this technique is that – as the teleport transition tended to have more extreme data points – filtering out extremely high error angles across all transition types introduces bias, and results in a reordering of mean values.

Counting the frequency of extreme data points would be one way to identify how often a participant became completely disoriented. An issue with this type of analysis is that the cutoff value is arbitrary, and could be selected such as to force a statistically significant result. As a result, we did not examine the frequency of extreme values.

In previous work that has used a pointing task, it was found that self-reported gaming experience played a role in pointing task performance [33]. In this work, data from gamers and non-gamers was separated, with the data from gamers having lower variance and producing a statistically significant result. A Spearman's rank-order test, however, found no correlation between frequency of playing video games and pointing task performance in our data. Participant conformity during the pen-and-paper map-placement task appeared highly variable, so this data was not analysed.

4.1.2 Time

As can be seen in Figure 8, there were outliers in the time taken to complete the pointing task. As a result, a Friedman test was run to determine if there were differences between transition types. The median time taken to complete the pointing task increased from teleport (Mdn = 8.453), to model (Mdn = 8.975), to Möbius (Mdn = 9.756), but the differences were not statistically significant, $\chi^2(2) = 1.613$, p = .446. There was little difference in the mean time to complete the pointing task between transition types, with teleport averaging 10.0s (SD = 4.8), model averaging 10.1s (SD = 3.7) and Möbius averaging 10.4s (SD = 3.7).

4.2 Subjective Measures

As subjective measures were given on a five point scale, it could be argued that the data was of an interval type. The data, however, was not normally distributed as assessed by Shapiro-Wilks, and contained outliers. As a result, parametric techniques were not appropriate. Therefore we treated the data as ordinal, and used the non-parametric Friedman test for analysis. Boxplots for all rated subjective metrics are shown in figure 9.

4.2.1 Moving Through the Space

A Friedman test was run to determine if there were differences in participants' subjective experience of moving through the space for the three different transition types. Participants' ratings for moving through the space were statistically significantly different for different transitions, $\chi^2(2) = 18.907$, $p < .0005$. Pairwise comparisons were performed

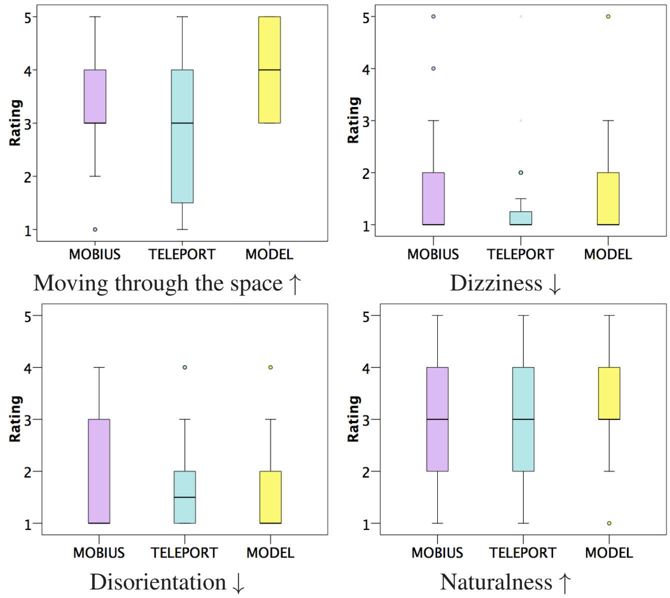

Fig. 9: Boxplots of subjective ratings. ↑ indicates a higher rating is better, ↓ indicates a lower rating is better.

using pairwise Wilcoxon signed-rank tests with a Bonferroni correction for multiple comparisons. Means are included here due to equal median values. Post hoc analysis revealed a statistically significant increase in the subjective experience of moving through the space from teleport (Mdn = 3.0, mean = 2.742) to model (Mdn = 4.0, mean = 3.875) ($p < .0005$) and teleport to Möbius (Mdn = 3.0, mean = 3.419) (p = .042), but not between model and Möbius (p = .063).

4.2.2 Dizzy

A Friedman test was run to determine if there were differences in participants' subjective experience of feeling dizzy during the three different transition types. Means are included here due to equal median values. Participants' ratings for feeling dizzy increased from teleport (Mdn = 1.0, mean = 1.389), to model (Mdn = 1.0, mean = 1.625), to Möbius (Mdn = 1.5, mean = 1.719), but the differences were not statistically significant, $\chi^2(2) = 2.774$, p = .250.

4.2.3 Disoriented

A Friedman test was run to determine if there were differences in participants' subjective experience of feeling disoriented by the three different transition types. Participants' ratings for feeling disoriented increased from model (Mdn = 1.0), to Möbius (Mdn = 1.25), to teleport (Mdn = 1.5), but the differences were not statistically significant, $\chi^2(2)$ = 2.136, p = .344.

Table 2: Transition preferences.

Winner	When playing against		
	Model	Teleport	Möbius
Model	-	35	46
Teleport	25	-	39
Möbius	14	20	-

4.2.4 Naturalness

A Friedman test was run to determine if there were differences in participants' subjective experience of naturalness during the three different transition types. Means are included here due to equal median values. Participants' ratings for naturalness decreased from model (Mdn = 3.0, mean = 3.317), to Möbius (Mdn = 3.0, mean = 2.952), to teleport (Mdn = 3.0, mean = 2.911), but the differences were not statistically significant, $\chi^2(2) = 2.482$, p = .289.

4.3 Preference

Binary preference data was analysed by fitting a Bradley-Terry model [8]. Here, a "contest" is considered to mean that a participant was asked to state a preference between two transitions, with the preferred transition becoming the "winner" of that contest. The count of wins for each transition type are shown in Table 2.

The parameters of the Bradley-Terry model were estimated using maximum likelihood. The model aims to estimate the probability that transition type i would beat transition type j in a contest, for each possible pairing of transition types i and j, where $i \neq j$. A positive-valued ability score α is calculated for each transition type, such that the odds that i will beat j are α_i/α_j.

The model can be expressed in the logit-linear form

$$logit[pr(i\ beats\ j)] = \lambda_i - \lambda_j,$$

where $\lambda_i = log\alpha_i$ for all i. This allows all of the parameters $\{\lambda_i\}$ to be estimated using standard generalised linear models (GLM).

Analysis was conducted using the BradleyTerry2 package for R [15]. As the parameters are relative rather than absolute, the 3D model parameter λ_{model} is set to zero as an identifying convention.

Preference counts for transition types decreased from model (wins = 82, $\lambda_{\text{model}} = 0$) to teleport (wins = 63, $\lambda_{\text{teleport}}$ = -0.3904) to Mobius (wins = 34, λ_{mobius} = -1.1180).

As the model can be expressed as a GLM, it is possible to calculate an analysis of deviance table and perform a chi-squared likelihood ratio test to obtain significance values.

Participant preferences were statistically significantly different for different transitions (GLM: $\chi^2(2) = 25.744$, $p < .0005$). Post hoc analysis was performed by fitting pairwise Bradley-Terry models, with a Bonferroni correction for multiple comparisons. Post hoc analyses revealed a statistically significant decrease in preference from model (wins = 82) to Möbius (wins = 34) ($p < .0005$) and teleport (wins = 63) to Möbius (p = .038), but not between model and teleport (p = .59).

4.4 Movement Profile

In order to explore the movement profile of the user following the different transition types, a time metric was examined following each transition during the returning phase (i.e. when a participant was returning to a known target location, and therefore not searching for a new target). The metric was the time between the completion of one transition and the initiation of the next transition. A boxplot of these results is shown in figure 10.

The average time data contained outliers, as assessed by visual inspection of the boxplots. As a result, a Friedman test was run to determine whether there were statistically significant differences in the average time before the next transition following each of the three transition types. The time before initiating the next transition was statistically significantly different for different transition types, $\chi^2(2) = 21.355$, $p < .0005$. Pairwise comparisons were performed using pairwise Wilcoxon signed-rank tests with a Bonferroni correction for multiple comparisons. Post hoc analysis revealed a statistically significant increase in the time before the next transition from Möbius (Mdn = 3.33) to teleport (Mdn = 5.1) ($p < .0015$) and from model (Mdn = 3.45) to teleport ($p < .0015$), but not from model to Möbius (p = 1)

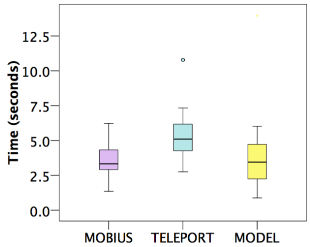

Fig. 10: Boxplot of average time before initiation of the next transition during the returning phase.

4.5 SSQ

The mean Total Severity (TS) score for the pre-experiment SSQ was 4.46 (SD = 7.81), while the mean TS for the post-experiment SSQ was 26.66 (SD = 32.13). TS values were calculated using the formula specified by Kennedy et al. [21]. While this is well below a high mean TS of around 70 [13], the increase from pre-exposure to post-exposure scores does indicate that simulator sickness may be an important factor in this context.

5 Discussion

5.1 Spatial Awareness

It is clear participants struggled with the pointing task, as is evidenced by the frequency of extremely high values in the error angles metric. The time taken to complete the pointing task was also high, with the mean for all three transition types being around 10s from fading to the grid environment to the participant pressing the button indicating they were pointing towards the target. This may indicate that MV360M experiences do not promote good spatial awareness. This is an important consideration, as an expected improvement in spatial awareness may be one of the most compelling reasons to employ such systems over single-view 360° media.

The poor spatial awareness results may also have been caused by the implementation of the experiment. As participants sat on a swivel chair, and could therefore not gauge their movement fully through proprioception, a loss of orientation may have been experienced. In two or three instances, participants had pushed themselves round and were rotating freely when the pointing task initiated, so had very poor orientation when the world faded to the grid environment. The grid environment should have ensured that participants retained visual cues about rotation, however, so the frequency of high error rates is perhaps still unexpected.

It is worthwhile to note that, while the procedure followed a pattern of pairs of targets as discussed in section 3.6, participants were not told of this pattern and in general did not appear to identify it. Each pointing task seemed to be unanticipated by participants. This means they may not have been making a specific effort at that time to maintain awareness of their own location or the location of the target. This could in part explain why the pointing task results are unexpectedly high.

There is also evidence that humans generally struggle with pointing tasks. In a study by Ruddle et al., a pointing task was used to assess spatial awareness in virtual buildings when viewed in a HMD or a desktop display [32]. In this study, participants navigated through a large-scale virtual environment using a keyboard and mouse in the desktop condition, and a handheld input device in the HMD condition. Based on visual inspection of their boxplots, the pointing task in the HMD condition produced average error angles of around 45°, while the average for the desktop display was around 55°. The standard error of mean (SEM) was approximately 5° for both. In comparison, our average error angles were between 41° and 51°, with a SEM between 4.5° and 6°. Their study, however, required participants to navigate buildings containing around 70 rooms – a task that would likely be considered substantially harder than ours. Other work in this area reveal similarly poor pointing task results (e.g. [31]). In retrospect, we believe our pointing task was not well designed for measuring spatial awareness. We would encourage future studies to either make the task less difficult, for example by not fading to a grid environment, or by choosing a different method for measuring spatial awareness.

5.2 Subjective Ratings

5.2.1 Disorientation

The subjective ratings for naturalness, disorientation and dizziness showed very little variation between transition types, with no differences nearing statistical significance. Each subjective question was asked once per transition type per participant. The question was randomly generated by the system, meaning the experimenter could not influence for which locations the question was asked. It was clear, however, that following occasional teleport transitions, some participants became disoriented. This usually presented itself through a verbal indicator from the participant. Anecdotally, this appeared to be more common when a teleportation ended close to a wall, meaning the participant had few visual features with which to orient themselves. Occasionally, participants would teleport twice in quick succession, resulting in confusion. The measurement method in use was not able capture these events. The metric does indicate, however, that the transition type did not generally have an impact on the subjective experience of disorientation.

5.2.2 Dizziness

The dizziness metric also did not establish any significant differences between transition types. Similarly to disorientation, there were occasional transitions where participants indicated verbally that they had felt dizzy, but our measurement method was not able to capture these. Anecdotally, these were during model transitions in which the locations were unusually far apart, meaning the user was moved faster to cover the larger distance over the 5s transition. Additionally, looking around during transitions seemed to increase dizziness.

While a simulator sickness questionnaire was administered before and after the experience, as all participants experienced all transitions roughly equally, it would not be possible to assess the cumulative simulator sickness effects of any individual transition type from our data.

5.2.3 Naturalness

During the Möbius transition, participants often commented on how it felt “weird”. It is perhaps surprising, then, that the naturalness metric did not detect any significant differences between transition types.

5.2.4 Moving Through the Space

It is clear from these results that the model and Möbius transitions created a stronger feeling of moving through the space than teleport transitions. This is not surprising for the model transition – as the user does virtually move through the space – but it is an interesting finding for the Möbius transition. The Möbius transition is image-based, and no additional information such as parallax is introduced. That the transition can induce the feeling of moving through the space means it could be a useful tool to allow the easy production of MV360M content that elicits this feeling, without the expense or complexity of reconstructing a 3D model of the scene.

5.2.5 Preference

As backed up by the quantitative preference results, the model and teleport transitions were generally well received. Participants expressed different opinions, with some preferring the teleport transition because it was faster, and some preferring the model transition because it was more fun or provided additional information about the scene. In general, participants tended not to enjoy the Möbius transition. Several participants reported finding it “weird”. One participant commented that it felt “like being pulled through a keyhole into a different space”. As stated earlier, however, our implementation is just one possible way to use the Möbius transformation for transitions. Other, more visually pleasing transitions may improve the user preference results. For example, in their work on the topic, Segerman et al. included the visual device of a picture frame to provide a join between two panoramas [14].

5.3 Movement Profile

As the teleport transition is instantaneous, it is perhaps not a surprise that following a teleportation the user takes more time to initiate the next transition than for the other two transition types. While we initially proposed that a longer delay before the next transition may indicate disorientation, during the study several other factors presented themselves as possible contributors to this effect. During the Möbius and model transitions, the user may have more time and visual information to decide on their next transition – saving them time on arrival. Additionally, the Möbius and model transitions give the user time to adjust their orientation during travel, allowing them to be facing in their desired direction on completion of the transition.

As the model and Möbius transitions each take 6 seconds to complete (0.5 seconds blend in, 5 seconds in transit, and 0.5 seconds blend out), teleporting would still be faster despite the increase in time before the next transition. It is interesting to note, however, that the Möbius and model transitions may not add as much total time to a journey as expected, as users take approximately 1.6 seconds longer on average following a teleport transition to initiate the next transition.

It is also interesting to note that there was very little difference between the Möbius and model transitions in terms of the delay before initiating the next transition. While the Möbius transition could feasibly have disoriented users, causing an increased delay, this does not appear to have been the case. Indeed, the median delay for the Möbius transition is slightly smaller than that of the model transition.

6 Limitations and Future Work

Although the SSQ was used to evaluate the simulator sickness effects of the experience on participants, it is not possible to identify which transitions contributed most to simulator sickness from our study design. Such an investigation would be valuable, as simulator sickness is likely to play an important role in the adoption of certain transition types.

As transition types were generally randomised, participants may have been unable to fully acclimatise to one transition type. Allowing a participant to acclimatise may be important to investigate aspects such as spatial awareness, as these may be affected by learning. Additionally, as participants acclimatise to the virtual space, their preference may change, with the faster teleport transition potentially becoming preferred over the more informative model transition.

In our MV360M content, the user explored scenes in which cameras were arranged in a connected network, with each camera being in the line-of-sight of at least one other camera. This allowed the entire scene to be explored, with available locations being visualised to the user by way of a sphere at that location’s position in the virtual space. This may not be the case for all MV360M content, as some scenes may be too sparsely captured for such a network to be feasible. Our research does not cover such scenarios, and the transition types explored may not be easily adaptable to these types of content.

Due to our desire to explore a 3D model transition, the available content was limited to static scenes. While our results may be applicable to dynamic scenes, there are other issues that were not addressed. As a result, further work is needed in order to understand the impact of dynamic MV360M on users.

There is a wide variety of transition types to be explored, including variations of the three transition types discussed here. For example, the Möbius and model transitions took six seconds each, irrespective of distance travelled. Varying the time based on the distance travelled could be one way to provide further information to the user. Additionally, our implementation of the Möbius transition is only one way to incorporate the Möbius transformation into a transition, and more complex transformations could potentially improve the visual appearance of image-based transitions. While there are many options, the methods in this paper highlight some important considerations, and show which metrics may be most sensitive to the transition type.

7 Conclusion

Our research investigates the impact of different transition types in MV360M for static scenes, in which users can navigate around a captured virtual space via a connected network of panoramic views. The three transition types explored were teleportation, a linear move through a 3D model of the scene, and an image-based Möbius transformation. The metrics investigated were spatial awareness, users' movement profiles, transition preference and the subjective feelings of moving through the space, disorientation, dizziness, and naturalness.

Our results indicate that the transition type has a significant impact on the subjective feeling of moving through the space, with the 3D model and Möbius transitions producing a stronger feeling of moving through the space than the teleport transition. The transition type also had a significant effect on a user's movement profile, with users taking on average 1.6 seconds longer to initiate the next transition following a teleport transition than a 3D model or Möbius transition. The subjective feelings of naturalness, disorientation and dizziness were not significantly different between transition types. A pointing task was unable to identify any significant difference in spatial awareness between transition types. These results indicate that the choice of transition type may have an impact on several aspects of the user's experience when exploring MV360M, and as a result content creators must think carefully before selecting a transition type.

Acknowledgments

This work was supported in part by grant EP/G037159/1 for University College London's Virtual Environments, Imaging & Visualisation Doctoral Training Centre (http://engdveiv.ucl.ac.uk/) from the UK Engineering and Physical Sciences Research Council (EPSRC) and BBC R&D. The authors would also like to thank Matterport for allowing us to use their content.

References

[1] D. Anguelov, C. Dulong, D. Filip, C. Frueh, S. Lafon, R. Lyon, A. Ogale, L. Vincent, and J. Weaver. Google street view: Capturing the world at street level. *Computer*, 43(6):32–38, 2010.

[2] D. N. Arnold and J. Rogness. Möbius transformations revealed. 2008.

[3] BigLook360. Use multiple camera views & 360 degree video to enhance live event broadcasts. http://biglook360.com/2011/11/360-degree-video-3/ [Online; accessed 08-September-2017].

[4] F. Bonato, A. Bubka, S. Palmisano, D. Phillip, and G. Moreno. Vection change exacerbates simulator sickness in virtual environments. *Presence: Teleoperators and Virtual Environments*, 17(3):283–292, 2008.

[5] D. A. Bowman, E. T. Davis, L. F. Hodges, and A. N. Badre. Maintaining spatial orientation during travel in an immersive virtual environment. *Presence: Teleoperators and Virtual Environments*, 8(6):618–631, 1999.

[6] D. A. Bowman, D. Koller, and L. F. Hodges. Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. In *Virtual Reality Annual International Symposium, 1997., IEEE 1997*, pp. 45–52. IEEE, 1997.

[7] E. Bozgeyikli, A. Raij, S. Katkoori, and R. Dubey. Point & teleport locomotion technique for virtual reality. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play*, pp. 205–216. ACM, 2016.

[8] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

[9] J. Carranza, C. Theobalt, M. A. Magnor, and H.-P. Seidel. Free-viewpoint video of human actors. In *ACM transactions on graphics (TOG)*, vol. 22, pp. 569–577. ACM, 2003.

[10] S. E. Chen. Quicktime VR: An image-based approach to virtual environment navigation. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pp. 29–38. ACM, 1995.

[11] A. Collet, M. Chuang, P. Sweeney, D. Gillett, D. Evseev, D. Calabrese, H. Hoppe, A. Kirk, and S. Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (TOG)*, 34(4):69, 2015.

[12] R. P. Darken and J. L. Sibert. A toolset for navigation in virtual environments. In *Proceedings of the 6th annual ACM symposium on User interface software and technology*, pp. 157–165. ACM, 1993.

[13] J. A. Ehrlich and E. M. Kolasinski. A comparison of sickness symptoms between dropout and finishing participants in virtual environment studies. In *Proc. Human Factors and Ergonomics Society Annual Meeting*, vol. 42, pp. 1466–1470. SAGE Publications Sage CA: Los Angeles, CA, 1998.

[14] eleVR. Spherical video editing effects with möbius transformations. http://elevr.com/spherical-video-editing-effects-with-mobius-transformations/ [Online; accessed 04-September-2017].

[15] D. Firth and H. L. Turner. Bradley-terry models in R: the bradleyterry2 package. *Journal of Statistical Software*, 48(9), 2012.

[16] J. Freiberg and R. Muntean. Bagan Four Buddha Temple. https://matterport.com/3d-space/bagan-four-buddha-temple/ [Online; accessed 01-December-2017].

[17] S. Freitag, D. Rausch, and T. Kuhlen. Reorientation in virtual environments using interactive portals. In *3D User Interfaces (3DUI), 2014 IEEE Symposium on*, pp. 119–122. IEEE, 2014.

[18] L. J. Hettinger, K. S. Berbaum, R. S. Kennedy, W. P. Dunlap, and M. D. Nolan. Vection and simulator sickness. *Military Psychology*, 2(3):171, 1990.

[19] Y. Horry, K.-I. Anjyo, and K. Arai. Tour into the picture: using a spidery mesh interface to make animation from a single image. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pp. 225–232. ACM Press/Addison-Wesley Publishing Co., 1997.

[20] D. M. Johnson. Introduction to and review of simulator sickness research. Technical report, Army Research Inst Field Unit Fort Rucker AL, 2005.

[21] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology*, 3(3):203–220, 1993.

[22] A. MacQuarrie. A 360° video of the Möbius effect. `https://www.youtube.com/watch?v=xV_hai4HUBU`, 2017.

[23] A. MacQuarrie. The effect of transition type in multi-view 360° media. `https://www.youtube.com/watch?v=XwdVenkQelY`, 2017.

[24] A. MacQuarrie. Spherical image editing. `https://github.com/andrewmacquarrie/spherical_image_editing`, 2017.

[25] L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pp. 39–46. ACM, 1995.

[26] Y. Morvan and C. O'Sullivan. Handling occluders in transitions from panoramic images: A perceptual study. *ACM Transactions on Applied Perception (TAP)*, 6(4):25, 2009.

[27] Oculus. Simulator Sickness. `https://developer.oculus.com/design/latest/concepts/bp_app_simulator_sickness/` [Online; accessed 04-December-2017].

[28] A. Okabe, K. Aoki, and W. Hamamoto. Distance and direction judgment in a large-scale natural environment: Effects of a slope and winding trail. *Environment and Behavior*, 18(6):755–772, 1986.

[29] Pitchfork. Watch Sigur Rós perform Kveikur songs live via interactive webcast. https://pitchfork.com/news/51227-watch-sigur-ros-perform-kveikur-songs-live-via-interactive-webcast/ [Online; accessed 08-September-2017].

[30] S. Razzaque, Z. Kohn, and M. C. Whitton. Redirected walking. In *Proc. EUROGRAPHICS*, vol. 9, pp. 105–106. Manchester, UK, 2001.

[31] A. E. Richardson, D. R. Montello, and M. Hegarty. Spatial knowledge acquisition from maps and from navigation in real and virtual environments. *Memory & cognition*, 27(4):741–750, 1999.

[32] R. A. Ruddle, S. J. Payne, and D. M. Jones. Navigating large-scale virtual environments: what differences occur between helmet-mounted and desk-top displays? *Presence: Teleoperators and Virtual Environments*, 8(2):157–168, 1999.

[33] S. P. Sargunam, K. R. Moghadam, M. Suhail, and E. D. Ragan. Guided head rotation and amplified head rotation: Evaluating semi-natural travel and viewing techniques in virtual reality. In *Virtual Reality (VR), 2017 IEEE*, pp. 19–28. IEEE, 2017.

[34] K. M. Stanney, R. R. Mourant, and R. S. Kennedy. Human factors issues in virtual environments: A review of the literature. *Presence: Teleoperators and Virtual Environments*, 7(4):327–351, 1998.

[35] P. W. Thorndyke and B. Hayes-Roth. Differences in spatial knowledge acquired from maps and navigation. *Cognitive psychology*, 14(4):560–589, 1982.

[36] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks Jr. Walking> walking-in-place> flying, in virtual environments. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pp. 359–364. ACM Press/Addison-Wesley Publishing Co., 1999.