Multi-Scale 7DOF View Adjustment

Isaac Cho, Jialei Li, and Zachary Wartell

Abstract—Multi-scale virtual environments contain geometric details ranging over several orders of magnitude and typically employ out-of-core rendering techniques. When displayed in virtual reality systems this entails using a 7 degree-of-freedom (DOF) view model where view scale is a separate 7th DOF in addition to 6DOF view pose. Dynamic adjustment of this and other view parameters become very important to usability. In this paper, we evaluate how two adjustment techniques interact with uni- and bi-manual 7 degree-of-freedom navigation in DesktopVR and a CAVE. The travel task has two stages, an initial targeted zoom and a detailed geometric inspection. The results show benefits of the auto-adjustments on completion time and stereo fusion issues, but only in certain circumstances. Peculiar view configuration examples show the difficulty of creating robust adjustment rules.

Index Terms—Multi-scale Virtual Environment, 3D User Interface, 7DOF User Interaction, Virtual Reality, View Adjustment

1 Introduction

Displayed on a typical desktop computer, interactive 3D environments require 6 degree-of-freedom (DOF) view control for 3D navigation; that is control of view translation and rotation (pose). However, this changes when the display and input hardware supports:

- stereoscopic display
- head-coupled display (typically via head-tracking)
- direct 3D manipulation (typically via 3D hand tracking or 6DOF input devices).

Define a "virtual reality" (VR) display system as a system that supports one or more of the above hardware components. Define a display system without any of these components as a "non-VR" system. The term "virtual environment" (VE) refers to any interactive 3D environment regardless of the type of display system used.

VR systems often include an additional view scale factor in the view model as a separate 7th DOF [1] [2]. The model is described by a coordinate system hierarchy containing the centers-of-projection (COP), the screen surfaces, and all 3D tracked input devices (6DOF) and user limbs. The model is best implemented using the same data types and APIs as the rendering engine's scene graph [3] [4]. The scale factor integrates into the computation of both in the viewing frustum(s) and the geometric virtual representation of any tracked 3D objects. Exactly one coordinate system contains the scale factor. This top level coordinate system is variously termed "room" [1], "platform" [5], "ViewPlatform" [3] "workspace" [4], etc. (This paper uses the term Workspace Coordinate System).

Therefore the "view scale" is the scale factor between the entire tracked physical interaction space and the virtual interaction space.

The view scale must be chosen with care for several reasons [6] [7]:

1) maximizing effective stereopsis and minimizing stereo fusion problems (as well as accounting for the eye's optical near point)
2) maximizing reachability during direct 3D manipulation and
3) optimizing head-coupled structure-from-motion cues

This paper refers to these items as the ***VR View Scale Issues*** (see [8] or [9] for a review).

In VR systems with a 7DOF view model, 3D UIs typically allow *dynamic* adjustment of view scale. This added DOF, however, complicates the navigation algorithm and user wayfinding. In contrast, non-VR systems typically implement a simpler 6DOF view model with 6DOF camera pose but without 7th view scale.

Related to the difference between 6DOF versus 7DOF view models, and VR versus non-VR systems is a third independent concept: the difference between a 'regular' VE and a multi-scale VE (MSVE). An MSVE is a VE that contains geometric details whose sizes cover several orders of magnitude [10] [9] [11].

Viewing the range of large and small details in an MSVE requires an "expansion" view maneuver. In 3D, this can be accomplished by either: dollying (view translation towards the target), scaling the view model about some chosen center-of-scale [1] [6], or zooming (narrowing the view field-of-view (FOV)). This paper uses the term view "expansion" to generically refer to all these specific view maneuvers.

With 6DOF view models, typical of non-VR systems, dollying is most common. For MSVEs a method is added to manually or automatically adjust the navigation speed [12], [13], [14], [15], [16].

In VR systems with 7DOF view models, both dollying or view scaling are possible. But for a VR system displaying an MSVE, dollying alone is rarely appropriate and view scale adjustment is also required to address the VR View Scale issues. Local terrain rendering systems with details down to meter precision [6] are sufficient to require a 7DOF view model to address the VR View Scale Issues. At a larger scale, early examples are the VR 3DUI's created for VGIS [17] [18], an out-of-core terrain rendering system that supports global terrain with details down to sub-meter precision, a precursor to Google Earth, etc. As an MSVE, VGIS require its VR UI to implement the 7DOF view model and use 7DOF navigation [10] [19] [20]. In all the interfaces, dynamic

- *Isaac Cho, Jialei Li and Zachary Wartell are with UNC Charlotte. E-mails: icho1, jli42, zwartell@uncc.edu*

Manuscript received April 19, 2005; revised September 17, 2014.

scale change is fundamental to navigation and interaction. A very broad range of applications, beyond geo-spatial, can leverage MSVEs [11] [21] [22]. Zhang et al. [21] tour a large range of potential scientific, medical and educational 3D applications that could benefit from viewing and interacting at dynamically changeable levels of view scale. They review the 7DOF view scale issues and examine the social implications for collaborative virtual environments when each user can individually dynamically change her view scale (called "avatar" scale in [21]).

Supporting 7DOFs for MSVEs on VR systems, raises several questions:

Q1: To what degree does semi-automated control of one or more of the 7DOFs improve user performance in traveling in an MSVE in terms of travel task completion time and management of VR View Scale issues? Further, how do different auto-adjustment algorithms compare to each other?

Q2: Do travel techniques that control all 7DOF with fewer modes (and hence have greater potential simultaneity of DOF control) improve travel task completion times?

Q3: Does splitting DOF control into more or fewer modes change the benefit of auto-adjustment?

This paper explores these questions. It examines the effect of three different view auto-adjustment conditions on a two phase travel task. The first phase (Task 1) involves a single, longer travel distance and requires less precise maneuvering. The second phase (Task 2) involves several shorter travel distances but requires more precise maneuvering (see Figure 3). Further, we compare two travel techniques. The one-hand (OH) travel technique has two switchable modes: mode 1 controls the 6DOF pose using the scene-in-hand [23] metaphor, while mode 2 controls the view scale using "hand-centered" scaling [1]. In this paper, we denote this use of 2 modes as ⟨6DOF+Scale⟩. The two-handed technique, called Spindle+Wheel (S+W) [24], has one mode which controls all 7DOF simultaneously. We denote this use of just 1 mode as ⟨7DOF⟩. Finally the experiment is run on two display systems: a DesktopVR system [25] (also "fish-tank VR") and a 3-screen CAVE [26].

Our prior work presents only the One-Hand interface [27] results. This paper contributes the following:

- Formal evaluation of both the two-handed and one-hand techniques with a comparison between them. This allows investigation of Q2 and Q3. Further, the two-hand experiment extends our understanding of Q1.
- Further details on the relation between scale change and travel time
- Elaboration of the "erroneous enlarge-and-push-away" problem, which seems to explain unexpected differences between the DesktopVR and CAVE results
- Detailed comparison of our experimental outcomes with prior work. This extends the understanding of Q1 and Q2.

2 Related Work

Head-coupled displays display 3D graphics where the generated perspective graphics image is dynamically adjusted based on head (or possibly eye pupil) position. Two types of head-coupled displays are Head-Mounted Displays (HMDs) and Head-Tracked Displays (HTDs). HMDs mount the displays on a headset or helmet [28]. In contrast, an HTD's [5] display is stationary mounted on a desk (DesktopVR [25]), a table (the responsive workbench [29]) or one or more walls (the CAVE [26]). "Mobile" HTDs exist as well [30].

We define stereoscopic 3D displays as a particular subclass of true 3D displays [31]. Stereo displays generate one or more pairs of optically planar images. For stereo displays, four eye separation values can be distinguished. First, the user's true eye separation (i.e. her interpupillary distance) may differ from the modeled eye separation, the value used in the view frustum geometry. Second these can be measured in either physical coordinates or virtual coordinates. The latter accounts for the 3D view (isotropic) scale factor [1]. Assume the human subject has a *physical* true separation of 6 cm while the *physical* modeled eye separation is set to 3 cm. Assume the 3D view scale used is $1/10^7$ so that a virtual Earth of approximate diameter 10^7 m is rendered at a diameter of 1 m. Then the *virtual* true separation is 600 km and *virtual* modeled separation is 300 km.

False eye separation occurs when the model and true eye separation differ. The modeled eye separation is deliberately set to a value other than the true value for purposes of distorting the depth of the presented stereo 3D image. False eye separation distorts the 3D stereo image by a non-affine homology [8], [10]. Importantly, this distortion is independent of the view scale value which controls the virtual/physical measurement ratios and induces a simple isotropic scale.

2.1 Stereoscopic Fusion Problems

Fusion problems in stereoscopic displays have been studied in stereo media [32] and computer graphics [33] and continue to be investigated [34] [35]. Barring dynamically adjusting the optical focal depth [31], for an (optically planar) stereoscopic 3D display fusion problems are managed in one of three ways: dynamic stereo depth range adjustment using geometric distortions, clipping out unfusible geometry, or simulating image blur from depth-of-field [36].

It is generally accepted that the perceived 3D image's depth range needs not to be equal to the modeled 3D depth range. Further ortho-stereo (view scale = 1 with no other distortion) is not necessary for many classes of applications.

Wartell et al. [8] classify 9 prior fusion control methods including perceived image depth adjustments and clipping plane methods, circa 2001. They abstract fusion control characteristics and match them to various application characteristics. Prior methods have a static or dynamic model of the near fusible distance (*nf*) and farthest fusible distance (*ff*) relative to the display screen. They compute a nearest point (*np*) and/or farthest scene point (*fp*) which are typically the nearest and farthest visible pixels. Dependent on the number of free parameters in a given adjustment technique, the adjustment can map *np* to *nf*, *fp* to *ff* or, map both simultaneously. Later Wartell [8] presents a 2 parameter method capable of mapping both.

Holliman et al. [37] present an adjustment technique which allows a defined region of interest in scene depth to have an improved perceived depth representation compared to other regions and which can keep this mapping constant even if total scene depth is changing. They also present a novel three-region algorithm for stereoscopic image adjustment.

Carvalho et al. [15] dynamically adjust stereo parameters based on a CubeMap structure [13] during the usage of two VR tools, fly and examine, in an MSVE.

2.2 MSVE Navigation

As mentioned, not all 3D UIs for MSVEs support 7^{th} DOF view scale because their underlying view model does not directly model it. This typical for rendering engines designed for non-VR systems. In these cases, view expansion (i.e. ”zooming in”) occurs through 6DOF view adjustment (dollying) with some auto-adjustment applied to travel velocity and possibly to the near/far clipping planes to manage z-buffer precision and clipping problems [1]. However, early VR systems demonstrate the importance of the 7DOF view model [1] and 7DOF travel techniques. Benefits for dynamically changing this view scale are demonstrated for travel techniques [6], object manipulation [7] and collaboration in multi-user VEs [38].

The following MSVE navigation review is partitioned between techniques based on 6DOF and 7DOF view models.

2.2.1 6DOF Navigation

Mackinlay et al. [12] develop the point-of-interest (POI) technique. With a mouse cursor the user designates a POI on an object. The system semi-automates orientation and flying speed as the user dolly’s in or out toward the POI.

Tan et al. [39] present a taxonomy for 3D navigation tasks and develop and evaluate the ”speed-coupled flying with orbiting” technique. It combines ego-centric and exo-centric navigation with automated speed and height-above-ground adjustment. The latter two aspects make it particularly relevant to MSVE’s. The display environment is desktop-and-mouse. Results suggest users perform better using their SCFO technique.

Wu et al. [40] present the design and evaluation of way-finding aids in an MSVE on a mouse-based non-VR system. The experiment compares 3 different way-finding aids: view-in-view map, animation guide and human system collaboration. The view-in-view map offers the best performance overall.

McCrae et al. [13] develop the CubeMap method that creates and samples a depth buffer in 6 view directions to modify travel speed and semi-automate orientation for MSVEs displayed on a non-VR, mouse-based system. Formal user studies are not reported.

Trindade et al. [16] improve two existing interfaces in order to assist and facilitate navigating an MSVE on a non-VR, mouse-based system. For an “automated” flying mode they include support for collision handling and automatic speed adjustment with respect to scale. For “automated” exo-centric travel, they extend a POI technique [12] with an automatic pivot point based on the construction and maintenance of a CubeMap. In a formal study of a non-timed exploratory task, users rate the automated UI version better than the manual one.

2.2.2 7DOF Navigation

As mentioned Robinett and Holliway introduce the concept of 7DOF navigation [1].

Ware et al. [6] present a z-buffer based auto-adjustment of both view scale (“cyclopean scale”) and the modeled/true eye separation ratio in stereoscopic VR systems. Later they support HTD’s [41]. They perform user studies to develop the auto cyclopean scale adjustment. They also test auto-adjustment effects on travel time [42]. All conditions use stereo display and automatic cyclopean scale without head-tracking. The 3D UI is mouse based and allows manual velocity control. The independent variable is whether a gain factor is also automatically applied to velocity based on the scene geometry. The users search a terrain for 3 boxes and identify the letters on the box sides. Compared to the no-gain condition, auto-gain reduces task time by roughly 30%. The experiment does not test the absence of all auto-adjustments, which is a condition in our present experiment.

Hanson and Wernert [43] present a general theoretical approach for constrained navigation in 6DOF and 7DOF environments using a CAVE example. Most of the given examples focus on constraining the pose 6DOF’s of the Workspace Coordinate System, but view scale is briefly discussed.

Wartell et al. [10] present an exo-centric, POI based 7DOF travel technique for VGIS, an out-of-core rendered, global terrain MSVE system on a responsive workbench. Auto-adjustment is a key part of the travel technique but no user studies are performed.

LaViola et al. [44] present “Step WIM” a clever hands free multi-scale navigation technique in a CAVE. Pierce and Pausch [45] design a travel technique for better scalability in large virtual worlds in an HMD system. Visible landmarks allow users to travel in the vicinity with a single gesture. In addition, symbolic place representations allow users to travel to distant locations with a small number of gestures. Houtgast et al. [19] (elaborated in Wartell et al. [20]) extend the VGIS responsive workbench UI for multi-scale volumetric weather visualization. They discuss balancing interaction, view scale, and stereoscopic display issues. They find a trade-off between direct manipulation and stereoscopic display, which must be optimized to help users perceive the environment. Their techniques rely on user created volumes-of-interest to drive auto-scaling. No formal study is performed.

Kopper et al. [11] design and evaluate 2 navigation techniques for MSVEs in an HMD system. A formal user study shows automatic scaling reduces travel task time compared to manual scaling for target-center based navigation. However, automatic scaling does not reduce time for a steering-based navigation. (More detail is reviewed in Section 8.5). Zhang et al. [22] use a 7DOF view model on a *non-VR* system – a mouse-based desktop. They compare 3 travel conditions: no-scaling, scaling-then-traveling and scaling-as-traveling (a generalization of travel using hand-centered scaling [1]). The travel techniques are essentially ego-centric and steering-based. All scaling is manual. In a formal study,

1. Scaling the view frustum(s) could be used to adjust the clipping planes. However, in a non-VR system the *exact same* visual result is also created by translating (dollying) the frustum and then adjusting the planes. On a VR system, these two alternatives are not equivalent regarding the VR View Scale Issues

contrary to expectations the scaling-as-traveling performed worse than no-scaling for both short and long distance targets. Scaling-then-traveling performed better than no-scaling for the long distance only.

Bacim et al. [14] design a framework for navigation techniques that provide understanding and classification of way-finding information (hierarchical and spatial information) needed for traveling in an MSVE in an HMD system. They compare four techniques: a steering+target travel technique based on Kopper et al. [11], Multi-Scale WIM [46], Hierarchically-Structured Map (HiSMap) – a new technique that presents a 2D tree view of hierarchy of nested LoS's, and MSWIM with HiSMap.

Oh and Hua [47] present a user study on 3 MSVE interfaces on a responsive workbench: focus + context, fixed focus + context, and overview + detail, with the purpose of identifying the differences of these interfaces with two tasks (path following and 3D map reading) in large scale information visualization on the 3D workbench.

3 User Interface

Before discussing our experimental design, we describe the user interface techniques that the experiment evaluates.

3.1 Navigation Techniques

The input devices are a pair of "buttonballs". They are 3.8 cm diameter plastic spheres containing a Polhemus Fastrak sensor and 3 buttons. Transparent spherical virtual 3D cursors represent each button ball. We apply an offset between the button ball and the spheres [48]. The UI accounts for the user's handedness by assigning button functionality based on the user's dominant hand. For brevity, further descriptions assume the user is right handed.

The one-handed technique (OH) works as follows. Holding one button engages a scene-in-hand technique [23]. Holding a second button engages rate controlled scaling where the center of scale is the cursor's position when the button is first pressed [1]. A separate, small red sphere shows this point while scaling is engaged. If the cursor is inside the Earth, we compute the intersection point of a line from the eye to the cursor with the Earth and display a small sphere at that point and this point is used as the center of scale instead [20]. A third button resets the view to handle getting lost.

Details on the two-handed travel technique, Spindle+Wheel (S+W), are discussed in [24]. It evolved from prior two-handed 6DOF work [49]. The key point is S+W uses a single mode (one button press). When engaged user hand motions can control all 7DOF. Prior work [24], shows S+W to perform faster for 7DOF docking tasks than the OH condition. For the experiment in this paper, a second button resets the view to handle getting lost.

3.2 Dynamic-View Adjustment Techniques

This paper evaluates two view auto-adjustments: auto-translation (AT) and auto-scale (AS). The auto-adjustments each only have 1 parameter (or DOF). Hence in any given circumstance, the AS and AT must choose to adjust for either the near or far stereo fusion limit.

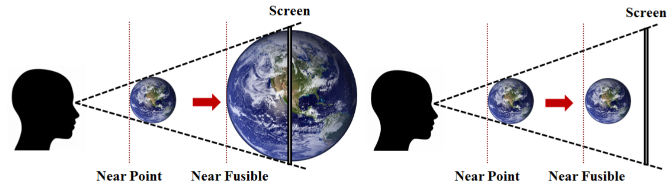

(a) View scale adjustment (b) View location adjustment

Fig. 1. Stereo adjustment techniques. (a) Auto-adjustment of view scale: the system pushes the Earth toward to the screen while preserving the retinal angle and projected image size. (b) Auto-adjustment of the view location: the system translates the Earth toward to the screen while preserving view scale.

AS uses a cyclopean scale [6]. For a near fusion violation, the cyclopean scale polls the z-buffer to determine the nearest scene point. If it extends outside the near fusible distance, the cyclopean scale is performed to make the near point fusible (transition indicated by the red arrow in Figure 1a). AT uses a view translation [10].

At each frame, an additional OpenSceneGraph render-to-texture pass renders to a "depth-buffer only" target at a quarter resolution of the display. The texture is transferred to the CPU which determines the nearest point by examining the entire depth image. If the nearest point is beyond the near fusible distance, the view translation is performed perpendicular to the scene (transition indicated by the red arrow in Figure 1b).

In order to control simultaneous violations of the near and far fusion points, AS and AT must be combined with an additional view parameter adjustment (Section 8.1). We chose to not incorporate a technique, because we argue the issue whether the view auto-adjustments interfere or aid 7DOF travel (Q1 and Q3) is best evaluated using auto-scale and auto-translation alone. Combining AS or AT with another technique would complicate experimental design and interpretation of results.

We use AT because while AT does perform some auto-adjustment, it does not alter the view scale, possibly avoiding interfering with the user control of scale. AT is roughly similar to Wartell et al.'s approach [10], but our travel technique is less constrained.

Figure 2 illustrates AT's complete algorithm. Near Target Distance (TD) is a static approximation of the nearest fusible distance. Far TD is a static, extremely conservative approximation of the far fusible distance. A z-buffer method determines the nearest point, *np*, in the scene. If $np < NearTD$ we translate perpendicular to the screen to bring *np* to *NearTD* (Figure 2a). If $np > FarTD$ we adjust to bring *np* to *FarTD* (Figure 2b). The auto-translation takes 0.5s to avoid abrupt stereo depth changes and user disorientation. If *np* is in the range [*NearTD*, *FarTD*], no adjustment occurs. This leads to a "buffer zone" such that if the nearest point is in the zone, no auto-adjustment occurs during user interaction.

For the CAVE, each of the three screens have a separate z-buffer. For auto-adjustment, only one camera's z-buffer is selected and polled, specifically the camera whose view frustum contains the Earth's center. Auto-stereo adjustment uses the TD's of the screen of the selected camera.

Since the CAVE has three displays for navigation, Shaw and Greens offset [48], perpendicular to the screen, must be modified. We implement a short-range, linear offset tech-

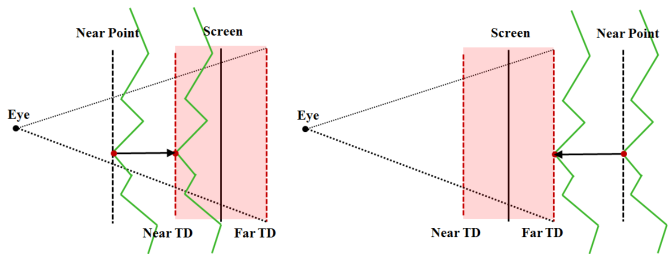

Fig. 2. Illustration of dynamic adjustment steps. (a) If the nearest point (*np*) is closer from the eye position than the Near TD, the system relocates the point to the Near TD. (b) If the *np* is further than the Far TD, the system relocates the point to the Far TD.

nique that supports a cursor offset in any direction (360°) similar on the Go-Go technique [50]. (Further details are found in the first author's dissertation [51] and [52]). The gain factor, however, is very low and allows a user 6′ from the center screen to extend the 3D cursor slightly beyond the screen. Importantly, walking forward a few feet with the arm at full extent allows the 3D cursor to reach the Far TD.

Once the user zooms in via view scale, far geometry can potentially exceed the far fusion limit. However, in a limited sense, the *Far TD* check does avoid far fusion problems due to the nearest scene point being too far away. Anecdotally we found this occurred frequently during 7DOF travel until we added the Far TD rule.

Our AS works similarly substituting the cyclopean scale for the translation. In both AS and AT, the geometry sampled from the z-buffer purposefully excludes the 3D cursors that represent the input devices. The 3D cursors are allowed to appear at any screen parallax. Three additional rules were added to address particularities of our using an extensive, out-of-core global terrain system [27]. Note, to avoid a possibly confounding effect of the z-buffer processing on frame rate, in the no adjustment condition the z-buffer is processed in a similar manner as in AT and AS but the information is not used. Because AS and AT do not include two auto-adjustable parameters, our conditions will generate more fusion issues than if they did incorporate a further parameter. In a deployed application an additional parameter (or perhaps gazed-tracked depth-of-field) would be added and we do not promote AS and AT complete fusion control solutions.

4 Experiment

Our experiment is designed based on the task of navigating a global, virtual Earth and visiting a place of interest and then inspecting details of the region. To motivate participants, we use pre-defined locations for the target destination at famous cities or landmarks in the world. In addition, we divided the world into spatial domains by its distance from the start position (America, Africa, Asia, Australia and Europe). This maintains similar travel distance across participants.

The completion time for traveling to the destination and inspecting the regions is recorded. The experiment examines the effect on completion time of the three different view auto-adjustment conditions (Section 3.2) across the two different navigation techniques (Section 3.1). Each experimental trial has two phases. In Task 1 the subject must find a target box. The box appears at one of four different sizes. This tests for any interaction of the auto-adjustment condition with the range of view scale change required to reach the target box. In Task 2 the subject must identify geometric details on a series of 4 new numbered boxes. These are the same size the original target box.

TABLE 1
Experimental Design

Subject	Factor	Values
Within	Task Type	Task1, Task2
Within	Box Size	1: 0.22%, 2: 2.2%, 3:22%, 4:220%
Within	Adjustment	NA: No Adjustment, AT: Auto-Translate, AS: Auto-Scale
Between	Interaction	OH: One-Handed, S+W: Spindle+Wheel
Between	Display	CAVE, Desktop VR
Between	Order	1: AT-NA-AS, 2: AS-AT-NA, 3: NA-AS-AT

The experimental factors are described in Table 1.

4.1 Hypotheses

Our primary hypotheses relate to the three questions (Q1-Q3) raised in the introduction:

H1: *For Task 1 and OH, AS and AT are strongly expected to have faster completion time than NA.*
Auto-adjustments reduce the DOFs the user must manually adjust. Task 1 requires 7DOF travel and OH in particular requires mode switching to control pose versus scale.

H2: *For Task 1 and S+W, AS and AT are expected to have shorter completion time than NA, but we expect adjustment benefits to be less than for Task 1 and OH.*
AS and AT reduce the DOFs the user must manually adjust and Task 1 requires 7DOF travel. However, unlike OH, S+W allows simultaneous control of all 7DOF (without mode switching), but potentially the AS and AT will help the user avoid navigating into view configurations that are uncomfortable to fuse.

H3: *For Task 1, AT is expected to have faster completion time than AS because AS may interfere with the user's desired manual scale.*

H4: *For Task 1 and 2, AS and AT are expected to produce fewer stereo fusion problems than NA.*

Further, we have three minor hypotheses:

H_i: *Spindle+Wheel will perform faster than OH.*
This is based on prior work [24].

H_{ii}: *For Task 1, the completion time should decrease linearly with the log of Box Size percentage:* Since Box Size 1 through 4 require varying scale change, completion time should decrease from 1 to 4. Further, since (1) the user interface provides proportional control of scale and (2) scale change is generally exponential, the completion time should decrease linearly with the log of Box Size percentage.

H_{iii}: *There will be an interaction effect between Task and Box Size on Completion Time:* The 4 numbered boxes in Task 2 always appear at roughly the same size as measured in physical screen space because Task 1 already performed the view scale change. Therefore, while Box Size should

(a) (b)

Fig. 3. Task 1 (a): The initial view is shown. Note the yellow wireframe docking box and the red arrow, a 3D cone, indicating the location of the target box which is 0.22% the size the docking box. Task 2 (b): After successful docking, four numbered boxes appear requiring detailed inspection.

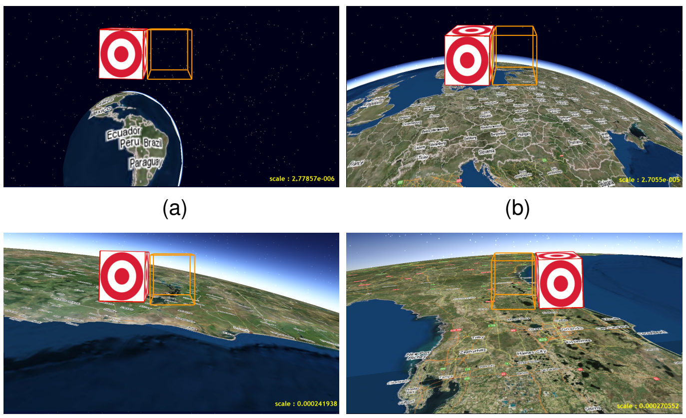

(a) (b) (c) (d)

Fig. 4. Four sizes of the target box, shown *after* the user has "zoomed in". Their sizes as percentage of the docking box are: (a) 220%, (b) 22%, (c) 2.2%, and (d) 0.22%. In virtual world coordinates their sides measure 5860 km, 1290 km, 28 km and 62.4 m.

have a strong effect in Task 1 Completion Time, H_{ii}, it will have a weaker effect on Task 2 Completion Time.

Finally we have one "expectation" [2]:

E1: *For Task 2, we expect AS and AT to have little or no effect as compared to NA.*
Reason: Unlike in Task 1, little or no view scale manipulation is required in Task 2. So AS's and AT's tendency to reduce the number of manually controlled DOFs is expected to be less beneficial. However, the AS and AT conditions may still provide some benefit in completion time by helping the user avoid navigating into view configurations that are uncomfortable to fuse.

4.2 Task

Each trial starts with global view of the entire Earth. In Task 1, there is a permanent, fixed sized wireframe "docking" box in the center of the display (Figure 3a). The user must travel to find the "target" box which may not be large enough to be initially visible as in Figure 3a. For tiny boxes, a red pointer hints at the target location. The user travels to bring the target box towards and then into a wireframe docking box. In Figure 4, the user has zoomed in and nearly completed the docking. As shown the target box can be one of four sizes which are 0.22%, 2.2%, 22%, and 220% of the wireframe box size. For shorthand, these target box sizes are named "Size 1" through "Size 4". A red pointer is needed for Size 1 and 2 and is displayed as a 3D cone. The cone's 3D size remains constant, as measured in screen space, during navigation.

After the user finishes Task 1, the target box and wireframe docking box disappear and four new, numbered boxes appear (Figure 3b). The new boxes are the same size as the target box from that trial. Each new box has a small hole on one face and a tiny colored sphere inside on the opposite face. The sphere color (red, blue or white) matches the colors of the button ball's buttons. The user must carefully maneuver to see the sphere color through the hole. The user indicates the sphere color by pressing the corresponding button on the left button ball. The user examines the boxes in order of their numeric labels (1 to 4). A success sound plays when the user presses the correct colored button. After the user presses the correct button for all four boxes, a new trial begins.

The user can reset the view position to the initial position by pressing a button of the right button ball during a trial. For Task 1, the initial position is where the user can see the entire virtual Earth (Figure 3a). For Task 2, the initial position is the position where the user finished Task 1.

At all times, a timer appears in the upper left of the screen. Below that is displayed the user's best time and then the current trial number. The upper right of the screen displays the auto-adjustment's engagement status as either "on" or "off". A name of the city, which is a target box location, is displayed below the auto-adjustment status. The view scale factor appears at the bottom right of the screen.

4.3 Participants

We recruited and ran all OH sessions first and all S+W sessions later. Overall, we recruited 48 participants (12 for each combination of display and interaction technique condition) from the Computer Science (CS) Department and the Psychology Department participant pool for the experiment. All participants have (corrected) 20/20 or higher eye vision.

The OH condition has 24 participants, 12 per display condition. In the DesktopVR group, 8 participants are CS major and 4 are non-CS major (8 are males and 4 are females). Participants have high daily computer usage (6.67, 1=*never* to 7=*a great deal*). 9 participants have experience with 3D UIs such as Microsoft Kinect. In the CAVE group, 6 participants are CS major and 6 are non-CS major (8 are males and 4 are females). Participants have high daily computer usage (6.42). 3 participants have experience with 3D UIs.

For S+W, there are also 24 participants. In the DesktopVR group, 7 participants are CS major and 5 are non-CS major (3 are males, and 9 are females). Participants have high daily computer usage (6.58). 6 participants have experience with 3D UIs such as Microsoft Kinect. In the CAVE group, 3 participants are CS major and 9 are non-CS major (7 are males and 5 are females). Participants have high daily computer usage (5.92). 6 participants have experience with 3D UIs.

4.4 Apparatus

The DesktopVR system uses a 24" Samsung 2233RZ display running at 120Hz at 1680×1050 resolution with Nvidia 3D

2. From an inferential statistics standpoint, we expect the null hypothesis to be true. However, inferential statistics tests can only reject the null hypothesis, so we label this item an "expectation"

Vision glasses. A Polhemus Fastrak tracks the glasses. The user is seated. A 3D cursor is displayed for each buttonball at a fixed offset, set by the user at start up. This allows the user to rest her elbows on the desk, her lap, or chair arm [48]. Based on informal pilot tests, the stereo TDs are ±8″ (20.32cm) from the screen.

The CAVE uses three Bacro Gemini^{TM} projected displays, 8′×6.4′ and 1280×1024 pixels each. Stereo is passive. A Polhemus Fastrak with a wide range emitter is used. The user stands with no place to rest her elbows or hands. The larger screen size causes the user to stand farther from the screen. This changes the fusible depth range in a non-linear fashion. For the CAVE, the TD is 48″ (1.21m) for the front screen and 36″ (0.91m) for left and right screens. During the pilot testing, we found that the user tends to stand approximately 6′ (1.52m) from the center screen and 4′ (1.21m) from right or left screen. When the user changes her view to the left or right, she tends not to move her body. We set shorter TDs, 38″ (0.96m) for left and right displays than the center one based on the observation that users tend not to physically walk left and right during the experimental task, but only forward and back. The software is built in C++ using osgEarth [53], VRPN [54] and osgVE [55].

4.5 Procedure

The participant first signs an informed consent form. For 10 minutes, she trains with a simple docking task to learn the travel technique. After training, the instructor teaches her about stereoscopic fusion problems by showing a case of extreme negative parallax. The instructor also explains how auto-adjustment can minimize fusion problems.

The participant then performs 20 trials for each auto-adjustment condition. In each trial, a target box appears in a random location with a random size for Task 1. Orientation of numbered boxes is also randomized per trial for Task 2. We record task completion time and number of resets for both Task 1 and Task 2. Adjustment order is counter-balanced between subjects using Latin squares.

5 Main Effect Results

We use the per-trial mean of task completion time and number of reset button presses. The reported F tests use α=.05 for significance. If Mauchly's Sphericity test fails, Greenhouse-Geisser adjusted F's are used and are denoted by appending ^{gg} to the variable name. Unless otherwise noted, planned post-hoc tests use Fisher's least significant differences (LSD) pairwise comparisons with α=.05 level for significance. Unplanned post-hoc tests use the Dunn-Sidak adjustment. The subjective preference data were analyzed by the Friedman test with α=.05. These post-hoc tests use Wilcoxon signed-rank tests.

Figure 5 shows completion times averaged over both OH and S+W. Table 2 shows significant main and interaction effects of the experimental factors.

Unsurprisingly, Task type has a strong main effect on completion time and Adjustment (H1,H2).

The interaction Display × Task is unexpectedly significant. Further, the simple effect of Display is also fairly strong. DesktopVR is faster than CAVE for Task 1 at 14.1s vs. 22.8s (p<.001) and Task 2 at 33.2s vs. 49.3s (p<.001).

TABLE 2
Significant effects of the experimental factors.

Factor	F	p	η_p^2
Overall			
Completion Time	F(1,26)=17.65	<.001	.949
Adjustment	F(2,72)=5.318	=.007	.130
Display × Task	F(1,36)=70.85	<.001	.329
Display	F(1,36)=67.85	<.001	.663
Display (Total Trial Time)	F(1,36)=61.18	<.001	.630
DesktopVR			
Task	F(1,18)=333.93	<.001	.949
Size	F(3,54)=41.86	<.001	.669
Travel Technique	F(1,18)=5.04	=.038	.219
Adjustment × Travel Technique	F(2,36)=9.43	<.001	.344
Adjustment (Total Trial Time)	F(2,36)=78.72	<.001	.814
DesktopVR - Task1			
Size	F(3,54)=82.42	<.001	.817
Travel Technique	F(1,18)=9.199	=.007	.338
Adjustment	F(2,36)=9.43	=.001	.344
Adjustment × Travel Technique	F(2,36)=3.91	=.029	.178
Adjustment (One-Handed)	F(2,18)=7.20	=.005	.444
Size (One-Handed)	F(3,27)=36.21	<.001	.801
Size (Spindle+Wheel)	F(3,27)=54.57	<.001	.858
DesktopVR - Task2			
Size	F(3,54)=8.38	<.001	.318
Adjustment × Travel Technique	F(2,36)=3.67	=.035	.170
Size (One-Handed)	F(2,36)=6.86	=.001	.433
Adjustment (Spindle+Wheel)	F(1,27,18)=11.49	=.004	.561
CAVE			
Task	F(1,18)=348.79	<.001	.951
Adjustment	F(2,36)=6.33	=.004	.260
Size	F(3,54)=108.72	<.001	.858
Task × Adjustment	F(2.36)=32.99	<.001	.647
Task × Size	F(3,36)=14.87	<.001	.452
Task × Adjustment × Size	F(6,108)=5.34	<.001	.229
Adjustment (Total Trial Time)	F(2,36)=50.73	<.001	.738
CAVE - Task1			
Adjustment	F(2,36)=3.52	=.04	.164
Size	F(3,54)=183.53	<.001	.911
Adjustment (One-Handed)	F(2,18)=6.95	=.006	.436
Size (Spindle+Wheel)	F(3,27)=97.23	<.001	.915
CAVE - Task2			
Adjustment	F(2,36)=16.78	<.001	.482
Size	F(3,54)=24.08	<.001	.572
Adjustment × Size	F(6,108)=3.55	=.003	.165
Size (One-Handed)	F(3,27)=8.33	<.001	.481
Adjustment (One-Handed)	F(2,18)=8.88	=.002	.497
Size (Spindle+Wheel)	F(3,27)=18.69	<.001	.675

This motivated analyzing Total Trial Time (which removes Task as a factor). Display is significant with DesktopVR at 54.2s and the CAVE at 81.1s. Further the average number of reset button presses per trial differs with DestopVR at 0.072 and CAVE at 0.322 (F(1,36)=4.96,p=0.03). Section 8.4 discusses the CAVE's lower performance. We believe a subtle design assumption interacted with a particularly error prone "cursor-centered scale view maneuver".

Hypothesis H_{ii} predicts in Task 1 a log-linear relation between Box Size and Completion Time. Due to the unexpected effect of Display, we examine the regression equation separately by Display using per subject average performance for each Adjustment × Hand combination. For DesktopVR, taking log_{10} of the Box Size percentage yields R^2=0.470 with m=-5.464 and b=18.748 ($F(1,286)$=253.3, p<.001). For CAVE, the results are: R^2=0.612, m=-8.933 and

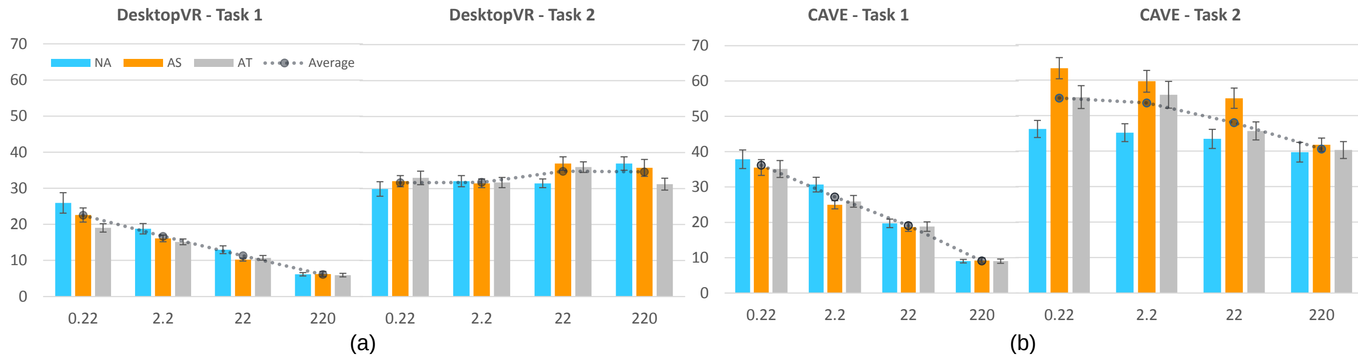

Fig. 5. Completion Time by Box Size for Conditions (averaged over Travel Technique). Error bars indicate standard error.

b=30.402 ($F(1,286)$=450.7, p<.001). Figure 5 labels the x-axis with percentage and uses equal tick spacing rather than plotting and labeling with $log_{10}(\%)$, but the linear trends are clearly visible. The predicted time for the 100% box sizes are approximately 12.5s for the CAVE and 7.2s for DesktopVR. This is roughly the time required for just 6DOF manipulations.

There is a significant Task × Box Size interaction as expected by H_{iii}. Repeating the earlier regression test for Task 2, finds that for DesktopVR the Box Size is not significant ($F(1,286)$=0.73, p=.393; R^2=0.003, m=0.425). But for the CAVE, the Box Size is significant ($F(1,286)$=40.22, p<.001) with R^2=0.123, m=-4.861 and b=53.36.

The "reset" button press count possibly indicates the user getting disoriented or an auto-adjustment performing a bad or unexpected maneuver. Over all trials, the average reset count is 0.393 per trial (contributed to by Task 1 and 2). But as mentioned, this differs for DesktopVR at 0.072 and CAVE at 0.322.

Further, Adjustment × Display is significant (F(2,72)=5.68, p=.005, η_p^2=.136). For CAVE, Adjustment is significant (F(2,36)=8.67,p=.001,η_p^2=.325). AS at 0.657 is greater than AT at 0.284 (p=.008) and NA at 0.24 (p.002) and AT is greater than NA (p=.036). For DesktopVR, Adjustment is significant (F(2,36)=8.20,p=.001,η_p^2=.313). AS at 0.137 is greater than both NA at 0.037 (p=.003) and AT at 0.041 (p=.001).

There are some statistically significant differences, but the per trial averages were less than 0.4 and not particularly illuminating.

Given the significance of Display and related interactions, we first split our analysis by Display type.

6 DesktopVR Results

Expected significant simple effects are Task, Size (H_{ii}) and Travel Technique (H_i). Not too surprisingly, the following 2-way interaction effects are also significant. Adjustment × Travel Technique is suggested by differing expectations for H1 vs. H2. Task × Adjustment is suggested by the difference in hypothesized effects in Task 1 (H1 and H2) versus the expected lack of effect in Task 2 E1.

The ratio of DesktopVR Task 2 to Task 1 time is 2.3. Given differing expectations of Adjustment effect on Task 1 and 2 (H1,H2), this ratio will effect Adjustment's effect on Total Time. Adjustment does have a significant effect on Total Time. NA and AS are 48.42s and 47.84s which are both significantly faster than AT at 66.23s (p<.001).

6.1 DesktopVR - Task 1

Figure 5a shows DesktopVR - Task 1. Size is significant as well as Travel Technique. S+W at 12.4s is faster than OH at 15.8s supporting H_i.

Adjustment is also significant (consistent with H1 and H2), and so is Travel Technique × Adjustment supporting the differing expectations for H1 and H2. Because Travel Technique is significant as well as its interactions, we split by Technique.

6.1.1 *One Handed and Spindle+Wheel*

For the OH technique, adjustment is significant as well as Size. LSD shows both AS, 14.8s, and AT, 13.7s, are faster than NA, 19.1s by 21.7% and 27.9% (p=.027, p=.012). This supports H1. AS and AT do differ significantly, failing to support H3.

Switching to the S+W technique, Size remains significant but Adjustment is no longer significant (p=.266). Section 7.2.2 will show similar results for CAVE Task 1-S+W.

6.2 DesktopVR - Task 2

Figure 5a shows DesktopVR - Task 2. Size is significant. This is unexpected (H_{iii}). The interaction Adjustment × Travel Technique is significant. We split by Travel Technique to see if Adjustment is significant for OH or S+W separately.

6.2.1 *One Handed and Spindle+Wheel*

For the OH technique, Size remains unexpectedly significant (H_{iii}). Adjustment is not significant.

But the effect size of Size for Task 2-OH (η_p^2=.433) is smaller than the effect size of Size for Task 1-OH (η_p^2=.800).

The overall lack of significance of adjustment is within expectations E1. If observed power were high, this could suggest that adjustment has no effect when the task requires little scale. However, observed power is low, 0.101. Hence, we look for replication of the non-significance of adjustment under S+W and the CAVE scenarios.

Unlike for OH, Adjustment is significant for S+W. Post-hoc test shows AT at 31.7s is slightly faster (6.4%) than AS at 33.84s (p=.029) supporting H3. NA and AT do not differ consistent with E1, but NA is faster than AS by 14.1% (p<.001), a reversal of E1. NA performing faster than AS suggests that in Task 2, where little scale change but much pose change is needed, the AS adjustment can be slightly harmful. For DesktopVR-Task 2-OH there was a non-significant trend in favor of AS and AT. But Section

7.2.2 will show further negative auto-adjustment effects for CAVE-Task 2.

6.3 Subjective Ratings

Users rated arm fatigue after the experiment for each adjustment condition (on a 7-point Likert scale, 1=*not at all* to 7=*very frequently*). The arm fatigue rate of participants is not significantly different by adjustment condition for the DesktopVR OH group (p=.135) nor for S+W group (p=.207).

Users rated the experience of stereo fusion problems on a 7-point Likert scale (1=*not at all* and 7=*very frequently*). For DesktopVR with OH, adjustment condition has no significant effect on a fusion problem rating for Task 1 (p=.798) nor Task 2 (p=.053). For S+W group, adjustment condition has no significant effect on the fusion problem rating for Task 1 (p=.558) nor Task 2 (p=.358).

Participants were asked if they preferred AS or AT to prevent stereo fusion problems and to perform tasks. In the OH group, 6 participants answered they prefer AS, 5 answered AT, and 1 had no preference. In the S+W group, 6 preferred AS, 4 preferred AT, and 2 had no preference.

7 CAVE Results

For the CAVE, the following simple effects are significant: Task, Adjustment, and Size. And several interactions are significant including: Task × Adjustment, Task × Size and Task × Adjustment × Size.

The CAVE average ratio of Task 2 to Task 1 time is 2.2. Hence given differing expectations of Adjustment effect on Task (H1,H2), the effect of Adjustment on Total Trial time will be strongly effected by this ratio.

Adjustment significantly effects Total Time. Completely contrary to expectations (H2) NA at 67.9s is faster than AT at 77.0s (p=.005) and AS at 98.5s (p<.001). NA is also faster than AS (p<.001).

The Task1/Task2 2.2 time ratio is an artifact that strongly contributes to this. We will show that it is with CAVE-Task 2, where auto-adjustments hurt performance. Figure 5b (Task 2) clearly shows this. Regarding the efficacy of auto-adjustment, we are most concerned with Task 1. Because Task itself and Task involved interactions are significant, we split analysis by Task.

7.1 CAVE - Task 1

Figure 5b shows CAVE - Task 1. The following simple effects are significant: Adjustment and Size. Unexpectedly, Travel Technique is not significant (p=.835) with S+W at 23.1s and OH at 22.7s contradicting H_{ii}. This is unlike DesktopVR - Task 1 (Section 6.1) and we suggest this is due to longer CAVE completion times which Section 8.4 will discuss. Even though Travel Technique is not significant, in order to simplify comparison to DesktopVR results, we split analysis by Travel Technique.

7.1.1 One Handed and Spindle+Wheel

For the OH technique, significant simple effects are Adjustment and Size. Regarding Adjustment, AS at 21.1s and AT at 21.7s are faster than NA at 25.3s (p=.006, p=.029). This supports H1 but not H3.

Overall, this suggests both CAVE and DesktopVR (Section 6.1.1), for Task 1 - OH, AS and AT perform better than NA, but in the CAVE the effect is less robust possibly muted by longer overall CAVE completion times.

Changing from OH to S+W, Adjustment loses significance completely. This was somewhat expected (H2). Size remains significant, see H_{ii}.

7.2 CAVE - Task 2

Figure 5b shows CAVE - Task 2. The following simple effects are significant: Adjustment, and Size. Sizes effect is unexpected (H_{iii}), but note its Task 2 (η_p^2=.572) is small compared to CAVE - Task 1 (η_p^2=.911).

Unlike with DesktopVR - Task 2, Travel Technique is just out of the significant range. This also occurred with Task 1 when switching between DesktopVR and CAVE. Section 8.4 argues this is due to longer CAVE completion times.

Further Adjustment × Size is significant. Figure 5B (Task 2) suggests this since for NA size has little effect, while for AT and AS size has an effect. The plot is visually quite different than DesktopVR - Task 2, Figure 5a, although DesktopVR's Adjustment × Size is just outside of significance.

For comparison between the CAVE and DesktopVR, we split analysis by Travel Technique despite non-significance, but first we consider Adjustment × Size.

7.2.1 Adjustment × Size

We split Task 2 by Adjustment to test for differential size effects. Confirming the visuals of Figure 5b (Task 2), for NA, Size is not significant (p=.06). Size is significant for AT ($F(3,54)$=12.02, p<.001, η_p^2=.400) and AS ($F(3,54)$=16.31, p<.001, η_p^2=.475). Pairwise unplanned comparisons (Dunn-Sidak) are as follows. For AT, Size 3 (45.7s) and Size 4 (40.2s) are both longer than each of Size 1 (55.2s, p=.016, p=.001) and Size 2 (55.9s, p=.03, p=.001). This is fairly visible in Figure 5b. For AS, Size 4 (41.8s) is faster than Size 1 (63.3s, p<.001), Size 2 (59.7s, p<.001), and Size 3 (54.9s, p=.003) and Size 2 is also faster than Size 3 (p=.043).

This behavior was very unexpected. Recall that Task 2, unlike Task 1, requires precise orientation control and little scale change to look into the holes in the four boxes. Size 1 through 3 boxes are very small relative to the Earth while Size 4 is roughly a quarter the Earth's size. We suspect that when trying to examine the smaller boxes, the auto-adjustment methods are interacting with the geometry of the scene (Earth + boxes) in such a way that the adjustments are making it more difficult to get the precise view orientation and location needed to look into the box holes. (A particular case of this is discussed in [24]). For the NA condition, this Adjustment × Size interaction is, of course, absent leading to relatively flat completion times for NA over Size as expected (H_{iii}). Section 8.3 discusses this in more detail. This problem is particular to the CAVE to be discussed in Section 8.4.

7.2.2 One Handed and Spindle+Wheel

For OH, Size is unexpectedly significant, see H_{iii}. Adjustment is also significant. Contrary to expectations (E1), NA at 44.3s and AT at 48.5s are both faster than AS at 56.3s, by 21.3% and 13.7% (p=.007, p=.025).

For S+W, Size is significant. This is unexpected because the numbered boxes are essentially always the same size, in physical screen space, at the start of Task 2. Adjustment is also significant. Contrary to expectations (E1) but consistent with OH (above), NA at 42.9s is faster than both AS at 53.6s (19.8%, p=.004) and AT at 49.9s (6.7%, p=.029). DesktopVR-Task 2-S+W also had a similar reversal of expectations (see Section 6.2.1).

7.3 Subjective Ratings

Users rated arm fatigue after finishing the experiment for each adjustment condition (on a 7-point Likert scale, 1=*not at all* to 7=*very frequently*). The arm fatigue rate of participants is not significantly different by adjustment condition for the CAVE group with OH (p=.093) nor with S+W (p=.638).

Users rated their experience of stereo fusion problems on a 7-point Likert scale (1=*not at all* and 7=*very frequently*).

For the CAVE with OH, Adjustment has a significant effect on fusion problems rating for Task 1 ($\chi^2(2)$=6.2, p=.045), but not for Task 2 (p=.115). For Task 1, median (IQR) fusion problem ranges for NA, AT and AS are [2,4];M=2, [1,2.75];M=2 and [1,2];M=2. Post hoc tests only show AS reduced fusion problems compared to NA (Z=.157, p=.031).

For the CAVE with S+W, for Task 1 adjustment condition has a significant effect on fusion problem ratings ($\chi^2(2)$=5.90, p=.048), but post hoc tests find no significant pair-wise comparisons. For Task 2, Adjustment has significant effect ($\chi^2(2)$=7.74,p=.019). Post hoc tests show only AS reduced fusion problems compared to NA (Z=-2.09, p=.02).

These fusion reduction results were weaker than expected. We expected both AT and AS to reduce the stereo fusion problems over NA for Task 1 and Task 2 (H4). Section 8.1 will suggest that our experiment design was not sensitive enough to detect differences.

In the CAVE with OH group, 3 participants preferred AS, 6 preferred AT, 3 had no preference and one disliked both. (Recall in the CAVE, AT and NA were faster than AS). In the S+W group, 4 participants preferred AS, 5 preferred AT, 3 had no preference. (Recall in the CAVE NA is faster than AT and AS).

8 Discussion

8.1 Stereo Fusion Control

We expected stronger and more consistent effects on reducing stereo fusion problems than we found (H3). However, this is perhaps not too surprising given that the auto-adjustment only manages negative-parallax limits and only indirectly addresses positive-parallax limits. Also, it may be too ambitious to expect users to rate fusion problems and report differences in a within subject design where the different adjustment conditions were presented back-to-back. A deployed version of any of the techniques covered in our experiment should employ an additional techniques to manage positive-parallax limits.

8.2 Task 1

For both DesktopVR and the CAVE, with the OH technique both AT and AS reduce completion time for Task 1, supporting H1 while under S+W there is no adjustment effect partially supporting H2. Averaging all Task 1 cases where AT or AS are significantly better than NA, the auto-adjustment techniques perform 20.3% faster. Overall AT and AS appear equivalent which does not support E1.

Prior work [24] compared Spindle+Wheel to the same One-Handed technique for a 7DOF docking task on a DesktopVR system. Their task is similar to Task 1, but additionally requires matching orientation alignment between the docking and target box. In [24], there was no auto-adjustment and the virtual environment included only the target box and a checkerboard ground plane. The experiment tested whether S+W, by allowing simultaneous 7DOF control, leads to faster completion times than OH, which requires scale/pose mode switching. This prior work used 3 target box sizes: 25%, 100% and 400%. S+W had faster completion times than OH for scaling cases 25% & 400%; the S+W averages 15.4s vs. OH at 19.0s.

The DesktopVR-Task 1 results are consistent with the prior study. In Task 1, which requires scale change, in Deskop VR, S+W is faster than OH, 11.7s versus S+W 14.1s (p=.017). But in Task 2, which requires minimal scale change, S+W performs the same as OH (30.4s vs. 33.3s, p=.229). Again, our CAVE results however show no difference for Task 1 OH vs. S+W, but we suspect this is due to the overall longer CAVE times. Overall, these results partially support H_i.

For both DesktopVR and the CAVE in Task 1, S+W loses a significant effect of Adjustment partially supporting H2. While the observed power for S+W is very low (0.06 and 0.266), the fact that the null hypothesis (H2) is not rejected in two replicated studies *between* subjects and across *different* display types (which have significant differences in overall completion time), suggests that switching from OH to S+W reduces the effect size of Adjustment. Further, there is a plausible explanation. Users can learn to manage the VR View Scale Issues reasonably well on their own with either OH or S+W. With OH mode, switching is required and auto-adjustment helps reduce this and hence travel time. However, for S+W has only one mode so adding auto-adjustment does not further reduce travel time.

Finally, for Task 1 completion time did decrease with Box Size as expected H_{ii}. There was the expected interaction effect between Task and Box Size H_{iii}. In detail, for Task 2 DesktopVR Box Size was not significant while for the CAVE Task 1 had a steeper slope (-8.933) than Task 2 (-4.861).

8.3 Task 2

Task 2 requires little scale change but it requires precise location and orientation control. Hence E1 expects AS and AT performance should be the same or possibly slightly better than NA. Results vary across combinations of OH and S+W and DesktopVR and the CAVE. In DesktopVR-OH, Adjustment had no effect (essentially supporting E1), but the other three cases Adjustment had the reverse of the desired effect – NA performed better than the auto-adjustment conditions. For all the cases with statistically significant differences, NA was on average 17.3% faster.

These results indicate that in certain view configurations the implemented auto-adjustments are harmful during tasks where scale change is minimal but precise pose control is

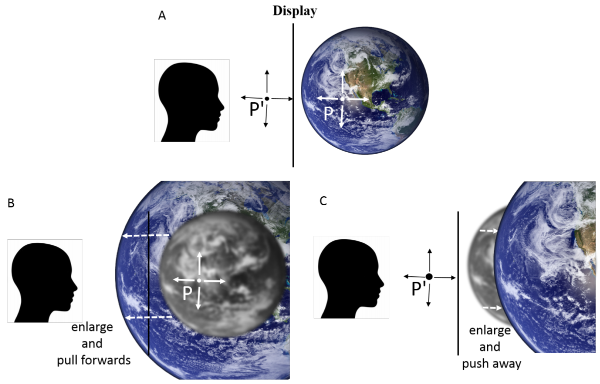

Fig. 6. Importance of cursor position during scale operation

needed. As discussed in Section 3, implementing precise auto-adjustments can be challenging. Since both the user and algorithm have control, it can be difficult to make an algorithm that always does the right thing at the right time. (Note, prior work that tests presence/absence of auto-adjustment in 7DOF travel has not separately analyzed maneuvers that require mostly scale change from those requiring mostly pose change). Our tested auto-adjustment techniques clearly need refinement. In future work, we want to track user input velocity. If the user is performing precise maneuvers, auto-adjustment is temporarily disabled.

8.4 CAVE - Far TD Problems

Across all conditions, CAVE completion times are longer than DesktopVR ones. Additionally, correlated with slower CAVE times, the overall total number of button presses per trial also increases (p=.013) with CAVE at 15.4s and DesktopVR at 12.6s. Also recall the reset count is DesktopVR at 0.072 but CAVE at 0.322. This implies that it is not merely larger or slower hand motions in the CAVE leading to longer completion times, but users are performing a larger number of maneuvers and likely getting disoriented

We hypothesize that the CAVE performs worse due to an increasing number of anecdotally observed, erroneous "enlarge-and-push-away maneuvers". This issue appears to be characteristic of both one-handed or two-handed scene-in-hand techniques where the center-of-scale is determined by a 6DOF cursor controlled via direct (isotonic) 6DOF input. First we describe the maneuver and then we describe why we suspect the CAVE interface increased its occurrence.

8.4.1 Erroneous Enlarge-and-Push-Away Maneuver

Exocentric travel techniques perform scale and rotations around a fixed point (or pivot point) that is not located at the view center of projection. The fixed point can be either 1) a POI which is a point selected lying on a visible surface [12] or 2) "cursor-centered" where the pivot point lies wherever the user has placed the 6DOF isotonic cursor. A third option is scaling-as-traveling where the 3D cursor can be arbitrarily positioned at any distance, but this option performs poorly [22]. Our tested techniques fall into the cursor-centered category. The cursor-centered, exocentric travel technique's ability to not have to designate a POI on a visual surface increases flexibility. However, there are use-cases (of desired view maneuvers) where being able to place the cursor close to a surface or inside a object, such as the Earth in our experiments, makes rotation and scale manipulations easier. User's subjective rating of the scene-in-hand metaphor for object rotations is known to be sensitive to the user's ability to place the pivot point inside an object which they are inspecting [56]. The scaling maneuvers under discussion are obviously related.

Consider the view configurations shown in Figure 6. Let $A{\to}B$ denote the view maneuver transitioning from view configuration A to B and $A{\to}C$ denote the maneuver transitioning between view configuration A to C. The points labeled P and P' are two possible centers-of-scale of a view scale maneuver. These would correspond to the cursor position in the OH technique or the Spindle center in the S+W technique. In either view maneuvers,$A{\to}B$ and $A{\to}C$, the scale factor is the same and the Earth visually enlarges. However, in $A{\to}B$ the Earth's surface moves towards the user while in $A{\to}C$ it moves away from the user. During our anecdotal observation, sometimes a user would desire to enlarge the Earth and pull it towards herself, $A{\to}B$, but instead she would enlarge and push the Earth away because the cursor was position P', and not P. Notably the authors (with 6, 2 and 20 years of experience with 7DOF travel in VR), rarely make this mistake presumably due to our understanding and experience with 7DOF travel. This error presumably occurs due to people's lack of experience of 7DOF view scale change in the physical-world. Logically in 7DOF travel techniques that are ego-centric, steering-based [11] or exo-centric point-of-interest based [10], erroneous enlarge-and-push-away maneuvers do not occur.

8.4.2 Relation with Far TD

The overall longer CAVE completion times and that fact that for the CAVE in Task 2 the auto-adjustment conditions were consistently worse than no adjustment may be caused by an increase in the number of erroneous enlarge-and-push-away maneuver's. There is reason to believe that the Far TD setting in the CAVE increased these occurrences.

In hindsight, we set stereo TD distances based on stereo fusion considerations (Figure 2) without giving enough attention to cursor range and reachability. Recall, in DesktopVR, the combination of the fixed translation offset and the Far TD location generally allowed one to easily place the cursor close to or inside the Earth (Figure 6b) at the Far TD. With the CAVE however, even with the non-linear offset, the cursor could reach the Far TD, but only if the user *walked* several feet towards the screen from the CAVE center.

We suspect that during our pilot testing and formative development, our users tended to walk forward/backward within the CAVE, but during our experiment, subjects tended to not walk much at all (possibly this could be an effect of fatigue). The latter behavior leaves the Far TD out of reach of the cursor. This would cause more frequent erroneous enlarge-and-push-away maneuvers, in particular under the AT and AS conditions which sometimes auto-adjust to the Far TD. (Recall, for the CAVE, Adjustment effects reset count with AS at 0.657 $>$ AT at 0.284 $>$ NA at 0.024). Further this would increase the difficulty of placing the cursor, also the rotation center, inside or near objects the user desires to rotate around for inspection, such as the 4 numbered boxes.

We draw several lessons about porting 3D UIs across different HTDs, such as DesktopVR and a CAVE. We present these from most specific to most general. Specific to only 7DOF auto-adjustment techniques that use a screen-distance based target distance, either:

- the target distance (our Far TD) should be placed at the most conservative distance implied by both stereo fusion and reachability considerations, or
- if the target distance is set based only on stereo fusion limits, the cursor offset gain must be set so that a user standing in her resting position (the middle of the CAVE) can still reach the target distance with the cursor

Slightly more generally, for any 7DOF travel technique using cursor-centered scaling, adding additional visual cues might help the user recognize when they are accidentally performing an enlarge-and-push-away maneuver (while intending to perform an "enlarge-and-pull-closer" maneuver). Possibly, if our participants recognized the nature of this error more readily in our CAVE condition, they would have walked closer to the screen. Also this visual feedback could be a useful training aid regarding accidental enlarge-and-push-away maneuvers.

Most generally, for any 3D cursor-based interactions on HTDs, cursor offsets and gain factors should account not just for the distances between the screen and nominal resting position, but also for whether users prefer to stand still or move within the space (for a CAVE, this is within the footprint of the CAVE). Moreover, designers should probably assume that a fatiguing user will start to walk less within a CAVE during long sessions and be sure to guard again any 3D UI aspects that are sensitive to the user's standing location. Finally, our results seem to indirectly support Li et al's [52] observations regarding short-range offsets in the CAVE.

8.5 Comparison to Prior Work

As reviewed in Section 2.2, we were not aware of any prior work that compares presence/absence of auto-adjustments on travel task completion time while partitioning the travel task into phases one of which is heavily scale changing and one of which is mostly view pose changing. Task 1 showed an overall 20.3% time reduction for auto-adjustments. Task 2 showed an overall 17.3% time increase for auto-adjustments. If the Task 1 and 2 completion times were balanced 50/50, the Total Trial time would have shown a modest 6.2% improvement for auto-adjustment.

We suspect, the Task 2 deficit could be removed by addressing the CAVE issue and re-testing for cases in Task 2 where auto-adjustment interferes with user goals. With a 50/50 balance and a 20.3% Task 1 improvement, Total Trial improvement could be up to 10.1%.

The above is speculative, in this experiment the presence/absence of an auto-adjustment only demonstrated a positive time effect in Task 1 on DesktopVR and CAVE for OH (H1). In other work, the presence/absence of an auto-adjustment had more marked effects. What may account for these differences?

First, we compare our results to Wartell et al. [10] which uses a 7DOF view model on a responsive workbench. The input device is a 6DOF virtual laser pointer. Travel is split into 3 modes: pan, a view scale centered on an user designated POI, and rotate. The travel technique maintains an exo-centric vantage point using several auto-adjustments. Wartell et al. state that in formative, informal evaluations, adding an auto-adjustment inspired by Ware et al. [6] made travel significantly easier. However, no formal user study was performed. Wartell et al. suggest merely disabling the auto-adjustment would generally cripple the travel technique because it was a integral part of its design. This is because the travel technique separated the DOFs into three different modes, none of which directly control view translation perpendicular to the screen. Disabling the auto-adjustment, would force the user to switch back and forth between the 3 modes to a unacceptable degree to avoid generating poor stereo images. In our experiment, however, the tested OH and S+W conditions allow translation in arbitrary directions. Wartell et al.'s [10] design goals were simply different since they aimed to mimic 2D pan-and-zoom map interfaces. More generally, we draw a tentative conclusion that travel techniques that use higher DOF input devices and have higher simultaneity of control (such as OH and S+W), will gain less completion time reduction from the auto-adjustment.

Next, Kopper et al. [11] use a 7DOF view model in a *monoscopic* HMD. The input device is a 6DOF joystick handle. They compare auto-scaling to manual-scaling in two travel techniques. The auto-scaling is managed using predefined volumes-of-interest, called Levels-of-Scale (LoS), displayed as bounding boxes around various 3D geometry. For both auto-scale and manual scale conditions, they evaluate two travel techniques, a steering-based technique (SB) and a target-centered based technique (TC). In SB, the user uses the 6DOF joystick to point in the direction of desired movement. Button 1 moves 'forward'; button 2 moves in reverse. In the manual scale sub-condition, the user presses a third button to scale up 10% and fourth to scale down 10%. In the LoS auto-scale sub-condition, the view scale is automatically changed when the user enters or exits a LoS box. In TC, the 6DOF wand controls a 3D cursor shaped like a magnifying glass. The object seen through the glass (along the line-of-sight) is highlighted. If a button is pressed, the user is smoothly and automatically translated to the LoS box. In the manual scaling sub-condition, the user presses one button to scale up 10% and another to scale down 10%. In the LoS auto-scale condition, in addition to the auto-translation, the view scale is smoothly interpolated to the view scale assigned to the destination LoS. Importantly, in TC when confined to a given LoS the only view translation and rotation possible is from physical head-motion. The HMD tracking system range appears to be short-range, only allowing one or two physical steps.

The travel task is more sophisticated and longer than ours and involves 2 data sets. With TC navigation, auto-scaling is far faster than manual scaling (Set 1: 49s vs. 126s, Set 2: 29s vs. 105s). With SB travel, however, auto and manual scaling are not significantly different (Set 1: 116s vs. 129s; Set 2: 77.21s vs. 118s).

Like our results, the presence/absence of auto-adjustment effectiveness depended on the interaction technique. However, enabling auto-scale for TC reduced travel time by 67%. This indicates that in the manual scale condition, TC users spent a far larger time adjusting scale than in SB. One interpretation is that TC's time efficiency is very sensitive to the user being at the right scale when inside a given LoS box and users have a difficult time manually adjusting to a 'good scale' to accomplish their task.

The LoS approach to auto-scale is excellent for use cases where a-priori LoS boxes can be defined. Because the auto-scale is not continually engaged once inside a new LoS, it cannot interfere with subsequent view maneuvers which our tested techniques do in Task 2. However, the LoS approach does not generalize to all use-cases, such as free exploration of systems like global terrain MSVEs nor MSVE volumetric data sets. The techniques evaluated in this paper do not assume an MSVE is populated by pre-defined LoS's. In such cases, to use LoS-based auto-scale approach the user would first have to *create* the "LoS" bounding boxes; this creation process must be added to the completion time of the travel task when comparing LoS auto-scale to other types of auto-adjustments. Houtgast et al. [19], [20] do something very similar to this in an MSVE with volumetric data, but present no formal evaluation. Our tested techniques both use the scene-in-hand metaphor while Kopper's use a steering technique (SB) and target-centered technique. In later work the authors [14] provide a hybrid SB-TC technique with auto-scale. Bacim et al. [14] show that 2D tree navigation through a hierarchy of the pre-defined LoS's is substantially faster than the hybrid SB-TC with auto-scale, but they also suggest interest in exploring exocentric techniques. The exocentric techniques we tested, with our suggested improvements, are an interesting starting point. Trade offs between egocentric and exocentric metaphors have long been observed [56].

9 Conclusion

Prior work shows that in VR systems used for MSVEs it often is necessary to use a 7DOF view model for managing the VR View Scale Issues. Further work shows that auto-adjusting some of the DOFs improves user performance but only under certain conditions.

This paper similarly found benefits of auto-adjustment versus manual control only under certain conditions. In DesktopVR and the CAVE, the results show a completion time advantage of 20% for auto-adjustment when the task requires scale change (Task 1) and when using a $\langle$6DOF+Scale$\rangle$ technique (OH). But the advantage disappears when using a (simultaneous) $\langle$7DOF$\rangle$ technique, Spindle+Wheel. For Task 2 (the inspection task requiring much less scale change) auto-adjustment was detrimental, but we believe we identified the CAVE specific problems and suggested further general refinements. The discussion sections described several lessons learned and several avenues for future work. Our long term goal is to develop and evaluate techniques for MSVEs with volumetric data and surfaces with complex bifurcations.

Acknowledgments

This work was supported in part by grant W911NF0910241 (PN 55836MA) from The U.S. Army Research Office.

References

[1] W. Robinett and R. Holloway, "Implementation of flying, scaling and grabbing in virtual worlds," in *Symposium On Interactive 3d Graphics*. ACM Press, 1992, pp. 189–192.

[2] ——, "The visual display transformation for virtual reality," *Presence: Teleoperators & Virtual Environments*, vol. 4, no. 1, pp. 1–23, 1995.

[3] H. Sowizral and M. Deering, "The java 3d api and virtual reality," vol. 19, no. 3, pp. 12 –15, may/jun 1999.

[4] G. D. Kessler, D. A. Bowman, and L. F. Hodges, "The simple virtual environment library: An extensible framework for building ve applications," *Presence: Teleoper. Virtual Environ.*, vol. 9, no. 2, pp. 187–208, 2000.

[5] D. A. Southard, "Viewing model for virtual environment displays," *J. Electron. Imaging*, vol. 4, no. 4, pp. 413–420, 1995.

[6] C. Ware, "Dynamic stereo displays," in *SIGCHI conference on Human factors in computing systems*, ser. CHI '95, New York, NY, USA, 1995, pp. 310–316.

[7] M. R. Mine, F. P. Brooks, Jr., and C. H. Sequin, "Moving objects in space: exploiting proprioception in virtual-environment interaction," in *annual conference on Computer graphics and interactive techniques*, ser. SIGGRAPH '97. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, pp. 19–26.

[8] Z. Wartell, "Stereoscopic head-tracked displays: Analysis and development of display algorithms," Ph.D. dissertation, Georgia Institute of Technology, August 2001.

[9] X. Zhang and G. Furnas, "mcves: Using cross-scale collaboration to support user interaction with multiscale structures," *Presence: Teleoper. Virtual Environ.*, vol. 14, no. 1, pp. 31–46, 2005.

[10] Z. Wartell, L. Hodges, and W. Ribarsky, "Third-person navigation of whole-planet terrain in a head-tracked stereoscopic environment," in *IEEE Virtual Reality*. IEEE Computer Society Press, March 1999, pp. 141–148.

[11] R. Kopper, T. Ni, D. A. Bowman, and M. Pinho, "Design and evaluation of navigation techniques for multiscale virtual environments," in *IEEE Virtual Reality Conference*. Alexandria, Virginia, USA: IEEE, March 25 - 29 2006, pp. 181–188.

[12] J. D. Mackinlay, S. K. Card, and G. G. Robertson, "Rapid controlled movement through a virtual 3D workspace," in *Annual conference on Computer graphics and interactive techniques*. ACM Press, 1990, pp. 171–176.

[13] J. McCrae, I. Mordatch, M. Glueck, and A. Khan, "Multiscale 3d navigation," in *Symposium on Interactive 3D Graphics and Games*. New York, NY, USA: ACM, 2009, pp. 7–14.

[14] F. Bacim, R. Kopper, and D. A. Bowman, "Design and evaluation of 3d selection techniques based on progressive refinement," *International Journal of Human-Computer Studies*, vol. 71, no. 7-8, pp. 785–802, July 2013.

[15] F. Carvalho, D. Trindade, P. Dam, A. Raposo, and P. Santos, "Dynamic adjustment of stereo parameters for virtual reality tools," in *Symposium on Virtual Reality (SVR), 2011 XIII*, 2011, pp. 66–72.

[16] D. R. Trindade and A. B. Raposo, "Improving 3d navigation in multiscale environments using cubemap-based techniques," in *ACM Symposium on Applied Computing*, ser. SAC '11. New York, NY, USA: ACM, 2011, pp. 1215–1221.

[17] P. Lindstrom, D. Koller, W. Ribarsky, L. F. Hodges, N. Faust, and G. A. Turner, "Real-time, continuous level of detail rendering of height fields," in *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, ser. SIGGRAPH '96. New York, NY, USA: ACM, 1996, pp. 109–118.

[18] P. Lindstrom, D. Koller, W. Ribarsky, L. F. Hodges, A. O. den Bosch, and N. Faust, "An integrated global gis and visual simulation system," Georgia Institute of Technology, Tech. Rep. GIT-GVU-97-07, March 1997.

[19] E. Houtgast, O. Pfeiffer, Z. Wartell, W. Ribarsky, and F. Post, "Navigation and interaction in a multi-scale stereoscopic environment," in *IEEE Virtual Reality 2005*, S. J. B. Frhlich and H. Takemura, Eds. IEEE Computer Society Press, March 2005, pp. 275–276.

[20] Z. Wartell, E. Houtgast, O. Pfeiffer, C. Shaw, W. Ribarsky, and F. Post, "Interaction volume management in a multi-scale virtual environment," in *Advances in Information and Intelligent Systems*, ser. Studies in Computational Intelligence, Z. Ras and W. Ribarsky, Eds. Springer Berlin Heidelberg, 2009, vol. 251, pp. 327–349.

[21] X. L. Zhang and G. W. Furnas, *Multiscale Space and Place*. Dordrecht: Springer Netherlands, 2005, pp. 261–280. [Online]. Available: http://dx.doi.org/10.1007/1-4020-3273-0_18

[22] X. Zhang, "Multiscale traveling: Crossing the boundary between space and scale," *Virtual Reality*, vol. 13, no. 2, pp. 101–115, 2009.

[23] C. Ware, "Using hand position for virtual object placement," *Vis. Comput.*, vol. 6, no. 5, pp. 245–253, 1990.

[24] I. Cho and Z. Wartell, "Evaluation of a bimanual simultaneous 7dof interaction technique in virtual environments," in *3D User Interfaces (3DUI), 2015 IEEE Symposium on*, March 2015, pp. 133–136.

[25] M. Deering, "High resolution virtual reality," *SIGGRAPH Comput. Graph.*, vol. 26, no. 2, pp. 195–202, Jul. 1992.

[26] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti, "Surround-screen projection-based virtual reality: the design and implementation of the cave," in *Annual Conference On Computer Graphics And Interactive Techniques*, ser. SIGGRAPH '93. New York, NY, USA: ACM, 1993, pp. 135–142.

[27] I. Cho, J. Li, and Z. Wartell, "Evaluating dynamic-adjustment of stereo view parameters in a multi-scale virtual environment," in *IEEE Symposium on 3D User Interfaces*, March 2014, pp. 91–98.

[28] I. E. Sutherland, "A head-mounted three dimensional display," in *Proceedings of AFIPS 68*, 1968, pp. 757–764.

[29] W. Krueger and B. Froehlich, "The responsive workbench [virtual work environment]," *IEEE Computer Graphics and Applications*, vol. 14, no. 3, pp. 12–15, May 1994.

[30] M. McKenna, "Interactive viewpoint control and three-dimensional operations," in *Symposium on Interactive 3D Graphics*. ACM Press, 1992, pp. 53–56.

[31] N. Holliman, N. Dodgson, G. Favalora, and L. Pockett, "Three-dimensional displays: A review and applications analysis," *IEEE Transactions on Broadcasting*, vol. 57, no. 2, pp. 362–371, 2011.

[32] L. Lipton, *Foundations of the Stereoscopic Cinema*. Van Nostrand Reinhold, 1982.

[33] N. Holliman, *Handbook of optoelectronics*. New York : Taylor & Francis, 2006, ch. Three-dimensional display systems, p. 10671099.

[34] M. T. M. Lambooij, W. A. IJsselsteijn, and I. Heynderickx, "Visual discomfort in stereoscopic displays: a review," in *Proc. SPIE, Stereoscopic Displays and Virtual Reality Systems XIV*, A. J. Woods, N. A. Dodgson, J. O. Merritt, M. T. Bolas, and I. E. McDowall, Eds., vol. 6490, March 2007, pp. 64 900I–1–64 900I–13.

[35] T. Shibata, J. Kim, D. M. Hoffman, and M. S. Banks, "Visual discomfort with stereo displays: effects of viewing distance and direction of vergence-accommodation conflict," in *Proc. SPIE: Stereoscopic Displays and Applications XXII*, vol. 7863, 2011, pp. 78 630P–78 630P–9.

[36] L. Leroy, P. Fuchs, and G. Moreau, "Real-time adaptive blur for reducing eye strain in stereoscopic displays," *ACM Trans. Appl. Percept.*, vol. 9, no. 2, pp. 9:1–9:18, Jun. 2012.

[37] N. Holliman, "Mapping perceived depth to regions of interest in stereoscopic images," in *In proceeds of Stereoscopic Displays and Applications XV*, San Jose, California, January 2004.

[38] J. Leigh, A. E. Johnson, C. A. Vasilakis, and T. A. DeFanti, "Multi-perspective collaborative design in persistent networked virtual environments," in *IEEE Virtual Reality Symposium*. IEEE, 1996, pp. 253–260,271–272.

[39] D. S. Tan, G. G. Robertson, and M. Czerwinski, "Exploring 3d navigation: Combining speed-coupled flying with orbiting," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '01. New York, NY, USA: ACM, 2001, pp. 418–425. [Online]. Available: http://doi.acm.org/10.1145/365024.365307

[40] A. Wu, W. Zhang, B. Hu, and X. Zhang, "Evaluation of wayfinding aids interface in virtual environment," in *Human-Computer Interaction. Interaction Platforms and Techniques*, J. A. Jacko, Ed. Springer Berlin Heidelberg, 2007, vol. 4551, pp. 700–709.

[41] C. Ware and D. Fleet, "Integrating flying and fish tank metaphors with cyclopean scale," in *Computer Graphics International*, Jun. 23-27 1997, pp. 39–46.

[42] ——, "Context sensitive flying interface," in *ACM Symposium on Interactive 3D graphics*. ACM Press, 1997, pp. 127–ff.

[43] A. J. Hanson, E. A. Wernert, and S. B. Hughes, "Constrained navigation environments," in *Scientific Visualization Conference (dagstuhl '97)*, June 1997, pp. 95–95.

[44] J. J. LaViola, Jr., D. A. Feliz, D. F. Keefe, and R. C. Zeleznik, "Hands-free multi-scale navigation in virtual environments," in *Proceedings of the 2001 symposium on Interactive 3D graphics*. ACM Press, 2001, pp. 9–15.

[45] J. Pierce and R. Pausch, "Navigation with place representations and visible landmarks," in *IEEE Virtual Reality Conference*, March 2004, pp. 173–288.

[46] C. A. Wingrave, Y. Haciahmetoglu, and D. A. Bowman, "Overcoming world in miniature limitations by a scaled and scrolling wim," in *IEEE Symposium on 3D User Interfaces*, March 2006, pp. 11–16.

[47] J.-Y. Oh and H. Hua, "Usability of multi-scale interfaces for 3d workbench displays," *Presence: Teleoperators and Virtual Environments*, vol. 17, no. 5, pp. 415–440, Oct 2008.

[48] C. Shaw and M. Green, "Two-handed polygonal surface design," in *ACM symposium on User interface software and technology*. ACM Press, 1994, pp. 205–212.

[49] D. P. Mapes and J. M. Moshell, "A two handed interface for object manipulation in virtual environments," *Presence: Teleoper. Virtual Environ.*, vol. 4, no. 4, pp. 403–416, 1995.

[50] I. Poupyrev, M. Billinghurst, S. Weghorst, and T. Ichikawa, "The go-go interaction technique: Non-linear mapping for direct manipulation in VR," in *ACM Symposium on User Interface Software and Technology*, 1996, pp. 79–80.

[51] I. Cho, "Bimanual stereoscopic interaction for 3d visualization," Ph.D. dissertation, Department of Computer Science, College of Computing and Informatics, University of North Carolina at Charlotte, August 2013.

[52] J. Li, I. Cho, and Z. Wartell, "Evaluation of 3d virtual cursor offset techniques for navigation tasks in a multi-display virtual environment," in *IEEE Symposium on 3D User Interfaces*, March 2015, pp. 59–66.

[53] gwaldron, "osgearth." [Online]. Available: http://osgearth.org/

[54] R. M. Taylor II, T. C. Hudson, A. Seeger, H. Weber, J. Juliano, and A. T. Helser, "Vrpn: a device-independent, network-transparent vr peripheral system," in *ACM symposium on Virtual reality software and technology*, ser. VRST '01. New York, NY, USA: ACM, 2001, pp. 55–61.

[55] E. Suma. (2010) osgvirtualenvironment. http://osgve.sourceforge.net/. University of North Carolina at Charlotte.

[56] C. Ware and S. Osborne, "Exploration and virtual camera control in virtual three dimensional environments," in *Symposium on Interactive 3D Graphics*, ser. I3D '90. New York, NY, USA: ACM, 1990, pp. 175–183.

Isaac Cho is a research associate at the Charlotte Visualization Center. His research interests include virtual reality, 3D user interface, 3D interaction technique, interactive visual analytics and human-computer interaction. He received his Ph.D. in computer science from University of North Carolina at Charlotte in 2013.

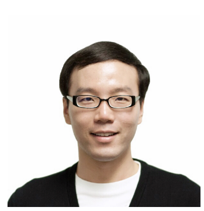

Jialei Li is a PhD candidate in the Department of Computer Science at the University of North Carolina at Charlotte. His research interests include 3D user interfaces, 3D interaction techniques and immersive virtual environments.

Zachary Wartell is a Associate Professor of Computer Science and Associate Director of the Charlotte Visualization Center at UNC Charlotte. His research interests are 3D user interfaces, stereoscopic and virtual reality displays, interactive 3D computer graphics and visualization.