

REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)	
	Technical Report	-	
4. TITLE AND SUBTITLE HyFinBall: A Two-Handed, Hybrid 2D/3D Desktop VR Interfacefor Visualization		5a. CONTRACT NUMBER W911NF-09-1-0241	5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER 611102	5d. PROJECT NUMBER
6. AUTHORS Isaac Cho, Xiaoyu Wang, Zachary Wartell		5e. TASK NUMBER	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of North Carolina - Charlotte 9201 University City Blvd. Charlotte, NC 28223 -0001		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211		10. SPONSOR/MONITOR'S ACRONYM(S) ARO 11. SPONSOR/MONITOR'S REPORT NUMBER(S) 55836-MA.6	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.			
14. ABSTRACT This paper presents the concept, working prototype and design space of a two-handed, hybrid spatial user interface for desktop VR targeted at users where a minimally immersive, desktop VR system is appropriate. The user interface supports dual button balls (6DOF isotropic controllers with multiple buttons) which automatically switch between 6DOF mode (xyz + roll,pitch,yaw), and planar-3DOF mode (xy + yaw) upon contacting the desktop. The mode switch automatically switches a button ball's visual representation between a 3D cursor and a mouse-like 2D			
15. SUBJECT TERMS visualization, virtual reality, user interfaces, geo-spatial, terrain			
16. SECURITY CLASSIFICATION OF: a. REPORT UU		17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 19a. NAME OF RESPONSIBLE PERSON Zachary Wartell 19b. TELEPHONE NUMBER 704-687-8442

Report Title

HyFinBall: A Two-Handed, Hybrid 2D/3D Desktop VR Interfacefor Visualization

ABSTRACT

This paper presents the concept, working prototype and design space of a two-handed, hybrid spatial user interface for desktop VR targeted at users where a minimally immersive, desktop VR system is appropriate. The user interface supports dual button balls (6DOF isotonic controllers with multiple buttons) which automatically switch between 6DOF mode (xyz + roll,pitch,yaw), and planar-3DOF mode (xy + yaw) upon contacting the desktop. The mode switch automatically switches a button ball's visual representation between a 3D cursor and a mouse-like 2D cursor while also switching the available user interaction techniques (ITs) between 3D and 2D IT's. Further, the small form factor of the button ball allows the user to engage in 2D multi-touch or 3D gestures without releasing and re-acquiring the device. We call the device and hybrid interface the HyFinBall interface which is an abbreviation for 'Hybrid Finger Ball.' We describe the user interface (hardware and software), the design space, as well as preliminary results of a formal user study. This is done in the context of a rich, visual analytics interface containing coordinated views with 2D and 3D visualizations and interactions.

HyFinBall: A Two-Handed, Hybrid 2D/3D Desktop VR Interface for Visualization

Isaac Cho Xiaoyu Wang Zachary Wartell
 icho1@uncc.edu xwang25@uncc.edu zwartell@uncc.edu
 UNC Charlotte

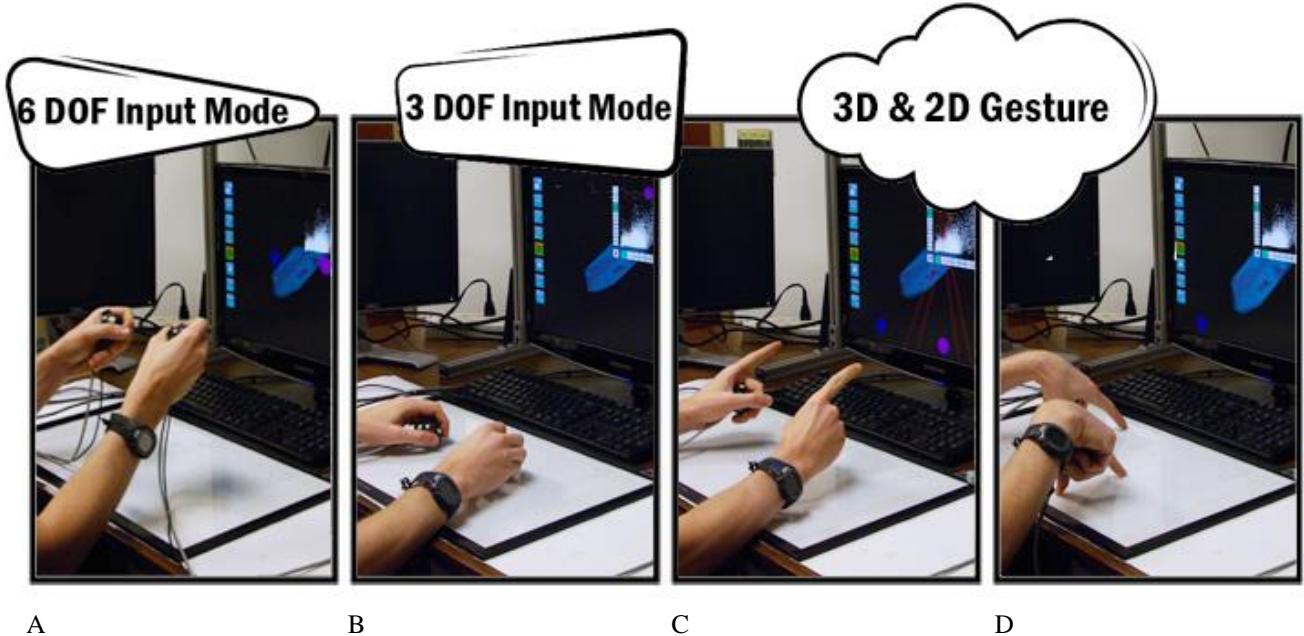


Figure 1: The HyFinBall UI supports 6DOF isotonic input (A), planar-3DOF input (B), 3D hand and finger tracking and gesture (C) and multi-touch (D). Note, the horizontal, multi-touch display is projected and disabled in this image, but see Figure 2 and the Supplemental video.

Abstract— This paper presents the concept, working prototype and design space of a two-handed, hybrid spatial user interface for desktop VR targeted at users where a minimally immersive, desktop VR system is appropriate. The user interface supports dual button balls (6DOF isotonic controllers with multiple buttons) which automatically switch between 6DOF mode (xyz + roll,pitch,yaw), and planar-3DOF mode (xy + yaw) upon contacting the desktop. The mode switch automatically switches a button ball's visual representation between a 3D cursor and a mouse-like 2D cursor while also switching the available user interaction techniques (ITs) between 3D and 2D IT's. Further, the small form factor of the button ball allows the user to engage in 2D multi-touch or 3D gestures without releasing and re-acquiring the device. We call the device and hybrid interface the HyFinBall interface which is an abbreviation for ‘Hybrid Finger Ball.’ We describe the user interface (hardware and software), the design space, as well as preliminary results of a formal user study. This is done in the context of a rich, visual analytics interface containing coordinated views with 2D and 3D visualizations and interactions.

Index Terms—stereoscopic display, virtual reality, user interface, two-handed interface, hybrid user interface, multi-touch, gesture, finger-tracking

1 INTRODUCTION

The ubiquitous Windows-Icon-Menu-Pointer (WIMP) user interface and its 2D mouse user interface techniques began with Xerox Parc’s and other’s seminal work. Similar to 2D interaction techniques (IT’s [1]), 3D ITs often require physical devices (e.g. ChordGloves [2](also pinch gloves), a bat [3], Cubic Mouse [4]) to provide a full six degrees of freedom (DOF) interaction to the user. Furthermore, HCI research has explored direct inputs by human modalities, such as voice, gaze, and gestures, for more natural ITs than those offered by physical input devices. Researchers have placed a particular emphasis on the study of natural human hand modalities like multi-touch direct input, and 3D hand gestures. These techniques allow

direct user interactions with minimal learning and make IT’s more fluid.

In this paper we present a minimally immersive, desktop VR [5] interface for a visual analytic application that provides two-handed bat (3D mouse) input, two-handed 2D mouse input, multi-touch and 3D gesture. The primary devices are two 6DOF button balls. We used these previously [6], borrowing from the bat, the FingerBall [7], and the button-enhanced bat [8]. This paper presents the HyFinBall (“hybrid-finger-ball”) user interface described below:

- **HyFinBall:** The HyFinBall interface starts with a pair 6DOF tracked balls with multiple buttons. Each ball is 4.5 cm in

diameter corresponding to a ping-pong ball. The software user interface that has the following properties. When a button ball is held in the air (Figure 1A), a 3D cursor displayed and 6DOF interactions are active. When a button ball is placed on the desktop, the UI automatically switches from to treating it as 6DOF isotonic device to treating it as a planar-3DOF input device (xy-position + yaw) and the 3D cursor is replaced by a 2D cursor in the plane of the screen. Each button ball independently switches between a 6DOF and planar-3DOF mode. During this switch, the user interface techniques available for the button ball switch from 3D IT's to 2D IT's. There is a translational offset between the physical location of the HyFinBall and its displayed 2D and 3D cursors. 6DOF mode uses an elbows-resting posture [8] while planar-3DOF mode uses a hands-resting posture. Strong consideration is given to stereoscopic display issues in the desktop VR environment when displaying the cursors. In particular, certain planar-3DOF IT's use projected 3D cursors.

- **HyFinBall + Finger-Tracking:** The HyFinBall is small enough to hold in a precision grasp [7] and small enough to be held with only the pinky, ring finger and palm in an average adult hand. This leaves the thumb, forefinger and (possibly) middle finger free. The free fingers can either:

- interact on a horizontal 2D multi-touch desktop display

OR

- perform 3-finger 3D interaction and gestures when in 6DOF mode.

Importantly, these 2D and 3D finger-tracking modes can be engaged *without* incurring an acquisition time penalty, i.e. the user does not drop and pick-up the button ball to engage and disengage these finger interaction modes.

The concept of using a single device that switches automatically between 6DOF mode and planar 3DOF mode, while not new (such the VideoMouse [9], and Logitech 2D/6D Mouse [10]) has not, to our knowledge, been integrated into any rich application that requires both 3D interaction and 2D interaction across coordinated views. The design space implied by the HyFinBall interface has not been explored with respect to desktop VR environments (in particular its stereoscopic 3D component) and this type of interface been not been studied for one-handed UIs, let alone two-handed UIs. To our knowledge, there has been no demonstration of a hybrid user interface (HUI) where the user uses a small form factor 6DOF held-device with a precision grip that can be continuously held while allowing the free fingers to engage in 2D multi-touch and/or 3D gesture interaction.

A user study is in progress focusing on the core HyFinBall concept comparing it to a mouse, the planar-3DOF-only mode and 6DOF-only mode across a variety of 2D and 3D combination tasks. In this paper, we present the HyFinBall and HyFinBall+Finger-Tracking concept and prototype (hardware+software). We present our anecdotal observations and describe the design space of the resulting hybrid interaction techniques. Finally, we present some preliminary findings of the aforementioned user study. This is done in the context of a rich, visual analytics interface containing coordinated views with 2D and 3D visualizations and with strong consideration of stereoscopic display issues in desktop VR. The supplemental video demonstrates the system.

2 BACKGROUND AND MOTIVATION FOR DESIGN

Many researchers have introduced 3D UI techniques for VEs. Bowman et al. [1] conducted many of the most recent, broad reviews

of 3D UIs and ITs and have reviewed and evaluated a number of 2D and 3D ITs. They also have identified specifications of ITs that will improve the usability of 3D interactions in real-world applications and have proposed guidelines for future ITs [11]. Liu et al. explored modern ITs for 3D desktop personal computers (PCs) [12]. A number of other articles also include review of physical input devices for 3D UIs [13] [14], and ITs for a large displays [15].

Bimanual interaction enriches interaction because humans often use two hands to accomplish tasks in the real world. A significant amount of research shows the advantages of bimanual interactions [16] [17] [18] based on Guiard's Kinetic Chain theory that classifies different categories of bimanual actions [19].

Several taxonomy's of spatial input technologies (hardware) [20] have been created as well as taxonomies of 3D spatial user interaction techniques (software) [21]. Here we use the following coarse categorization of spatial input hardware:

- 2D vs 3D input
- held-devices vs body-tracking

Our operational definitions are as follows. A 2D input device only tracks within a physical plane. 3D input tracks motion applied in 3-dimensions (at least 3DOF position and up to 6DOF). *Held-devices* are spatial input devices held by the user, while *body-tracking* tracks the body (such as hands and fingers). Body-tracking *never* requires the user to grasp a prop, but it may require some encumbering mechanism (gloves, fiducial markers, etc.).

A traditional mouse is a 2D held-device with 2 position DOFs. A 2D mouse with the ability to yaw perpendicular to the motion plane [21] is referred to here as a planar-3DOF device. Multi-touch is a body-tracking, 2D input with roughly 20 DOFs (10 fingers x 2 position DOFs). VIDEOPLACE was a early body-tracked 2D interface [22]. Notably the user was completely unencumbered (i.e. requiring no worn apparatus of any kind, not even fiducial markers).

3D input interacts in a 3D space. The bat [3] is an isotonic, 3D held-device with 6DOF pose (position and orientation). A bending-sensing data glove with a 6DOF tracker attached is a body-tracking 3D input, not held-device input. The ideal implementation of body-tracking, of course, is a completely unencumbered system. Wang et al [23] demonstrate unencumbered hand+finger-tracking. The operational definition of body-tracking treats encumbered and unencumbered implementations as sub-categories.

Various researchers have demonstrated [7] [24] [25] that having a 3D held-device grasped in the hand is beneficial due to the tactile feedback (passive haptics) it provides for 3D manipulation. Such feedback does not exist in hand or finger-tracked 3D UIs, but does exist in 2D multi-touch UI's or haptic augmented 3D systems.

When considering a held input device, devices are held in either a precision grasp or power grasp. For some applications, such as a VR system for training a user to use a real-world tool, a power grasped prop is ideal—assuming the real-world tool requires a power-grasp. However, a precision grip allows finer control due the larger “bandwidth of the fingers”. Physically the HyFinBall device follows Zhai et al's FingerBall which had a single button activated by squeezing [7]. But the HyFinBall interface uses multiple buttons and is two-handed following Shaw and Green [8]. (We use these button balls in XXXXX et al [6] but that system does not contain any of the HyFinBall hybrid UI concepts). As a general purpose input device for desktop VR applications, we follow the above authors and promote using a pair of generic shaped devices that remain in the user's hands for relatively long durations to minimize device acquisition time penalties. This is opposed to using multiple, specially shaped 3D held-devices that must be put down and picked up repeatedly. We suggest that for data visualization applications (as opposed to VR training applications) a pair of generic devices (or perhaps a few devices of different but generic shapes [26]) will be superior for many application domains.

Early tangible user interfaces [27] were 2D held-devices that were planar-3DOF. Tangible interfaces were unique in that the user

had a multitude of different held-devices available on a horizontal display surface and the held-devices were untethered and required no power (an external camera tracks their 2D pose).

Most user interface devices and corresponding user interface techniques that provide spatial manipulation use either held-devices or body-tracking, but not both. There are some exceptions. For example, the touch mouse contains a multi-touch surface on the top of the mouse [28]. However, to our knowledge there has been relatively little development and experimentation with user interfaces that support 2D and 3D held-devices while simultaneously enabling 2D/3D hand+finger tracking. This is the goal of the HyFinBall+Finger-Tracking interface.

Ideally the HyFinBall button ball would be untethered allowing full 360 degree rotations without an encumbering, entangling cord. Bradley and Roth demonstrate untethered computer vision tracking of a fist-sized ball, but occlusion remains a problem, especially for a two-handed scenario. Current battery and sensor technology still precludes constructing an accurate, small-form factor wireless 6DOF ball, but this area of engineering is very active [29]. Finally, non-isomorphic rotation techniques [1] can ameliorate cord entanglement during rotation operations.

Mapes and Moshel [2] use an HMD with 6DOF tracked pinch gloves and a physical surface at a 45 degree angle. A pair of 3D cursors are positioned roughly corresponding to the position of the user's hands. When the hands rest on the surface they are supported and the pair of pinch-gloves essentially act like a pair of 3 button mice. However, the display of the 3D cursors remains the same regardless of hand position. In contrast, in the HyFinBall planar-3DOF mode, if the user rests the button ball on the desk it changes both the cursor display and the interaction techniques available. This difference is motivated in part, due to the display system difference, i.e. HMD in Mapes and Moshel vs desktop VR here. In the HyFinBall planar-3DOF mode, the 2D cursors are within the plane of the vertical display screen while the button balls remain on the desktop surface. This is designed specifically to mimic mouse usage and to place the 2D cursors at zero-screen parallax to simplify stereo viewing issues when interacting with the 2D GUI elements.

The term hybrid user interface (HUI) refers to a UI with multiple methods for spatial input, frequently supporting both bimanual or unimanual interaction and 2D and 3D interaction. Benko et al. [30] combine a multi-touch 2D surface with hand and finger 3D gestures and 3D interaction in an augmented reality system. They coin the terms HUI and cross-dimensional gestures.

Some earlier devices support a similar notion of cross-dimensional interaction. The VideoMouse [9] and the Logitech 2D/6D Mouse [10] are a single device that support both 6DOF mode and planar-3DOF mode. However, in neither system was this concept extensively developed into a hybrid 2D/3D UI nor was two-handed interaction supported. The utility of confining the motion of 6DOF device to a physical plane, such as a held tablet, to reduce the physically manipulated DOF's has been demonstrated [1]. However, these prior works do not use a significant displacement between the physical device and it's representative 2D or 3D cursor (as in [8]) and neither of these works UI's implement the 6DOF to planar-3DOF mode switching found in the HyFinBall interface.

Massink et al. introduced HyNet, an HUI system for desktop-based navigation [31]. This work uses a traditional mouse for navigating the 3D world with a conventional desktop system. However, the system only uses 2D GUIs with 2D UIs and does not provide a solution for 3D visualizations and VE systems. The authors also introduce a programming abstraction for the HUI, with traditional desktop-based systems that used conventional mouse and keyboard inputs. The HUI addresses both theoretical abstraction and 3D input modalities.

Alencar et al. present HybridDesk that combines 2D and 3D interactions with a tracked Wiimote and WIMP interface for an oil platform visualization [32]. There are three UIs in HybridDesk used to evaluate their HUI techniques: VR-Nav for navigation and selection, VR-Manip for manipulation, and the traditional WIMP UI.

More recently, Magic Desk [28] utilizes multi-touch input, a mouse, and a keyboard within a traditional desktop environment for unimanual and bimanual interactions. The authors explore suitable physical positions of multi-touch input relative to the user during the experiment. Althoff et al. present a multimodal interface for navigation in arbitrary virtual VRML worlds [33], which uses a mouse, keyboard, joystick, and multi-touch input. However, their environment was limited to 2D visualizations and 2D interactions. The Slice WIM interface, which uses a multi-touch table with a head-tracked, stereoscopic wall screen display for a medical imaging volumetric dataset [45], allows multi-touch interaction on the table to control 3D data using two widgets.

Multimodal user interfaces (MUI) generally use more than just spatial input; for instance they combine voice and gesture [34] [35]. Bolt introduces a system called "put-that-there," which uses voice and gaze inputs [36]. Within GIS systems, voice and gaze inputs also are popular interaction methods in MUIs [37] [38]. The main advantage of natural human input modes is that they do not require any held-device and users need less training.

HUIs and MUIs can be combined with augmented reality as well. ICARE is an example of such a mixed environment [39]. Bianchi et al. developed a hybrid AR system, which used a hybrid external optical tracker for the user's head pose and a subsequent visual landmark-based refinement of the pose estimation [40] that uses AR's overlaying of virtual objects on the user's real environment [41]. Other previous works include medical volumetric datasets designed for use by surgeons [42] [43].

Many HUI and MUI systems incorporate hand-held, mobile devices. Song et al. introduce an application called what-you-see-is-what-you-feel that uses a mobile device for input and a wall-mounted display for medical imaging volumetric data visualization [44]. Users employ 2D multi-touch input on the handheld device to manipulate the 3D medical volume data on the large wall-mounted display through the wireless network.

Researchers also can use HUIs and MUIs in collaborative systems. Each user can handle a different system employing heterogeneous displays with various techniques to share the visualization or data with other colleagues. Schmalstieg et al. introduced a mixed reality environment that combined AR, ubiquitous computing, and a desktop metaphor for a collaborative system used with medical volume data [42].

3 THE HYFINBALL USER INTERFACE



Figure 2: HyFinBall UI: Head-tracked stereoscopic vertical display, projected multi-touch table using PQLab's frame, dual button balls, and dual Kinects for 3D hand and finger-tracking.

We developed the HyFinBall interface in the context of an application for analysing terrain meshes from 10 years of LIDAR scans of the North Carolina Coast from the NOAA Digital Coast database. We refer to this application as DIEM-VR. In the first section we describe DIEM-VR 2D and 3D interactions, the visualization components and their coordinated view mechanisms. In the second and third sections we discuss how the HyFinBall interface is used to interact with these components. In the final sections we discuss how the HyFinBall+Finger-Tracking is currently implemented.

3.1 DIEM-VR: A Desktop VR System for Terrain Analysis

Our ultimate goal is to integrate our prior work on terrain change detection algorithms [46] into the DIEM-VR application. At present, we have focused on implementing features in DIEM-VR that motivate the HyFinBall interface. While the user can view any one of the 10 years of LIDAR terrain scans, we have not ported our change detection algorithms into DIEM-VR.

The user sits at dual screen, desktop VR system. It uses Nvidia 3D vision glasses and a Polhemus Fastak for head-tracking and for tracking the HyFinBall devices. Two Windows Kinects view the desk space running 3Gears finger-tracking software and a PQLab screen is place on the horizontal screen with an overhead projector.

The system displays a single patch of terrain which can be optionally color-coded by height or displayed as a wireframe mesh or point-cloud. A series of 2D menu buttons appears on the left of the primary screen. These implement a horizontal, pull-“right” menu. All 2D menu items are displayed at zero screen parallax.

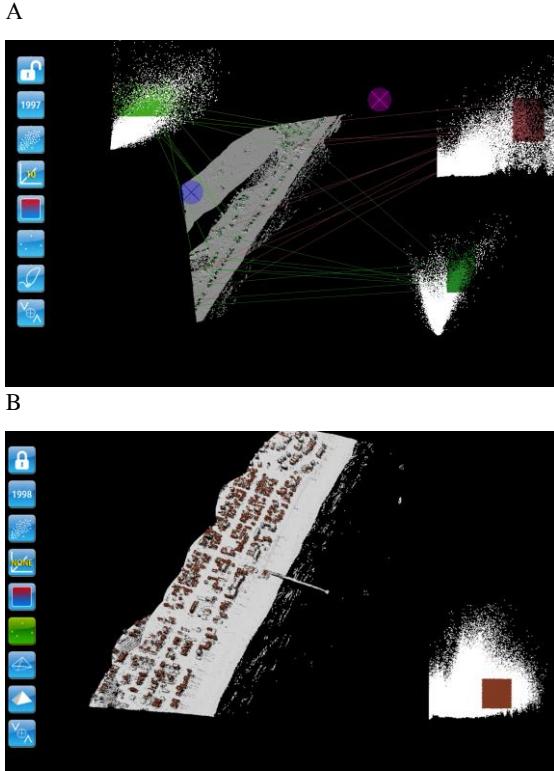


Figure 3: A) Point-cloud rendering of terrain patch and interactive, coordinated scatter-plot representations of LIDAR points B) Selection of LIDAR points in scatter-plot high-lights house roofs.

The user can add and delete multiple scatter-plots whose plot points each correspond to a terrain point. Each plot point's x-y location is determined by a geometric characteristic of the associated

terrain point such as the terrain point's average local slope, local degree of roughness, etc. In other words, each original terrain point has several additional geometric characteristics associated with it and by creating scatter-plots along these dimensions, the user can view the terrain in a different feature space such as plotting local roughness versus elevation. The scatter-plots are constrained to the zero-parallax plane. They can be repositioned manually or automatically. When a cursor hovers over a scatter-plot boundary, icons along the x or y axes appear allowing selection of the statistic that will be plotted on the given axis. Various statistics such as average gradient, maximum gradient, local standard deviation can be selected.

The user can brush points in the scatter-plot. Brushing occurs by creating a rectangular selection region. The selected points are highlighted on the terrain surface using a color pre-assigned to the scatter-plot. The user can optionally enable the display of lines connecting the scatter-plot points and the terrain points. This gives a stronger visual impression of how the brushed scatter-plot points are spatially distributed on the terrain. (For performance, only a randomly chosen subset of the connecting lines is drawn). Understanding the spatial structure of this “line net” is greatly enhanced by the stereoscopic display. It has some conceptual similarities with traditional 2D parallel coordinates. Figure 3A, shows three scatter-plots with line nets connecting their brushed regions to the terrain points. In Figure 3B, the scatter-plot in the lower-left plots elevation versus local gradient. The brown selection region is selecting for relatively low elevations with minimal gradient. This causes mostly house roofs to be highlighted in the terrain view.

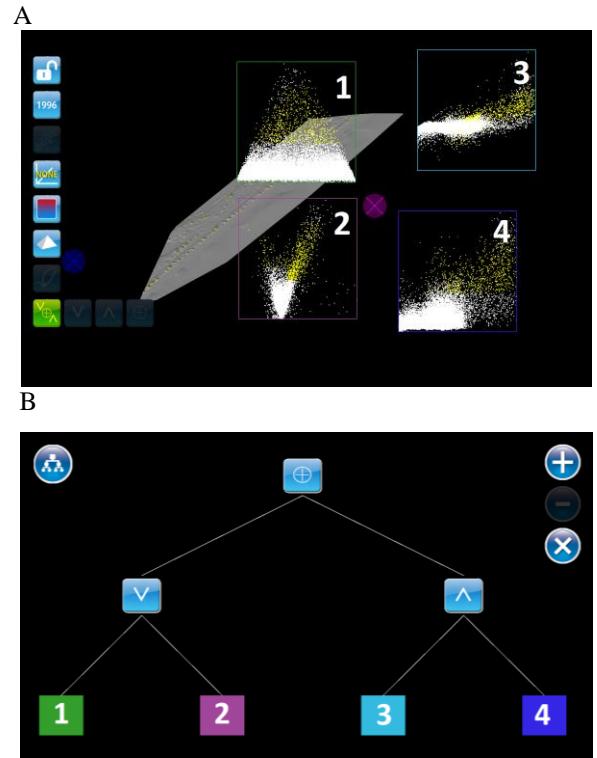


Figure 4: Scatter-plots with selected regions and interactive, Boolean expression tree.

After creating multiple scatter plots and brushing different regions in each scatter plot, the user can construct a Boolean expression that combines the different selections. Only the terrain points that satisfy the Boolean expression are highlighted in the terrain view. The horizontal multi-touch display shows the tree structure of the Boolean expression. Figure 4B shows a logical

expression of (1 OR 2) XOR (3 AND 4). Numeric labels map elements of the expression to the scatter plot. After saving the expression, an icon appears at left top of to record the expression. Users can delete, select or modify prior saved expressions.

The user can select 3D terrain points directly. LIDAR scans have multiple returns and are hence multi-planar (not strict height-fields). There are situations where one may want to select points not only within a certain foot-print but also within a limited height range. For example, the user might want to select tree top returns and not the lower layer returns from the underlying ground. While selection in these situations is not as complicated as selection within true volumetric data [48], we provide both a general 3D selection box interface and a 2D lasso selection interface for the cases where the user does not need to restrict the selection height range. This general capability for volume selection will be necessary when integrating true volumetric data into the terrain systems as we did in [48]. The 3D selection box can be created, moved, rotated and resized. This IT is detailed in Section 5.

The 2D lasso selection interface is optimized for height-field terrain and cases where the user wants to select only by specifying a footprint on the terrain. The user creates 2D lasso selection in the plane of the screen. This lasso selection shape is projected onto the terrain by using the 2D cursor location and projecting a ray from the cyclopean COP through the cursor position on the frustum projection window.

Finally, there is an individual terrain triangle selection mode. In this mode the terrain triangle underneath the 2D cursor is selected and all other terrain triangles within a range of similar height values are also selected. As the 2D cursor is dragged this selection is continuously highlighted. (Other criteria for selecting ‘similar’ terrain polygons are, of course, possible).

All these terrain region selections and scatter-plot selections use brushing-and-linking across these coordinated views that are updated in real-time.

These rich 2D and 3D interactions drive our core HyFinBall interface (6DOF + planar-3DOF modes) and the HyFinBall+Finger-Tracking interface.

3.2 HyFinBall – The Interface

In this section, we discuss the HyFinBall interface which employs the 6DOF mode, planar-3DOF mode and the automatic, independent mode switching of each HyFinBall when it contacts the desk surface. A user study is in progress comparing the following device conditions and UIs:

- I. the auto-switching HyFinBall UI
- II. dual planar-3DOF mode only UI
- III. dual 6DOF mode only UI
- IV. a single mouse

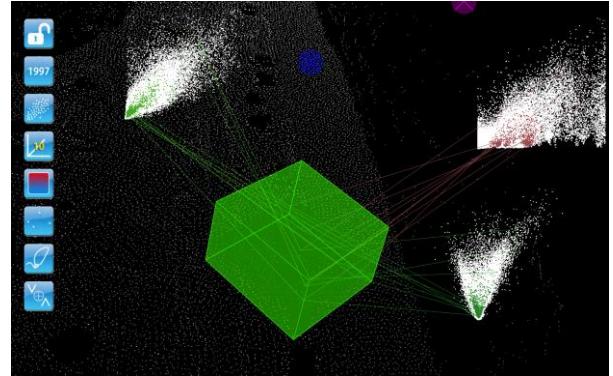
This comparison is done across a variety of 2D and 3D tasks in different sequential combinations.

In it’s 6DOF mode, the left HyFinBall implements a scene-in-hand metaphor [3] for camera pose manipulation plus separate 3D cursor centered view scaling [47]. In 6DOF mode the HyFinBall’s virtual representation is a transparent, blue sphere with a user adjustable translational offset [8]. When the left HyFinBall is placed on the desk, planar-3DOF mode is enabled. Now, the HyFinBall’s cursor is replaced by a transparent, 2D blue disc that always remains at zero screen parallax. This cursor interacts like a standard 2D mouse cursor for selecting the menu bar on the left. From our anecdotal observation and several pilot study participants, in the stereo display the switch from the 3D sphere cursor to the 2D disc cursor is immediately apparent.

The right HyFinBall’s 6DOF mode implements and initiates 3D selection box creation. In 6DOF mode the right HyFinBall’s virtual representation is a transparent, orange sphere with a user adjust translational offset. As described in 3.1, the selection box is used to select points on the 3D terrain. In details, the selection box creation

is a combination of the two-handed technique of Ulinski et al. [6] and a 3D widget [11].

A



B

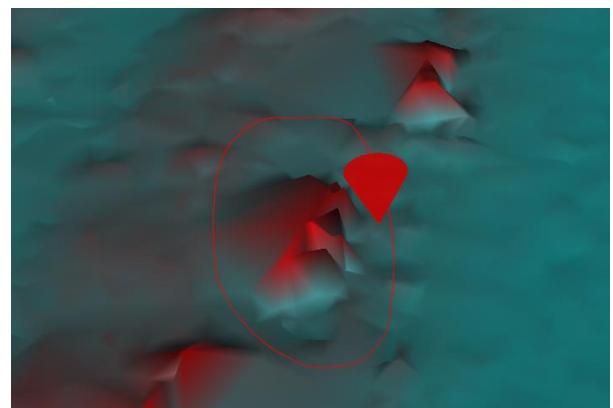


Figure 5: points selection (a) Selection box and (b) lasso

When the right HyFinBall is placed on the desk, the 3D cursor is replaced by a transparent, 2D orange disc that remains at zero screen parallax. In this mode, the orange disc acts like a 2D mouse cursor for interacting with any created scatter-plots. The user can move the plot or switch the plot data axis using button Button A. The user can select a rectangular region of scatter plot points with button Button B. In it’s planar-3DOF mode, the right HyFinBall can also be used for 2D lasso selection of the terrain points. In this mode, the orange disc is replaced by a different 3D cursor whose 3D position is the intersection of a ray cast from the cyclopean eye through the 2D cursor’s computed position on the frustum projection window. (This cursor may appear at negative or positive stereo parallax depending on the position of the intersected terrain point). In prior work, we used a similar technique where we replaced the display of the desktop 2D mouse cursor with projected 3D cursor. This enabled a mouse controlled travel technique option in our exo-centric, travel technique on stereoscopic virtual workbench [48].

3.3 HyFinBall - Design Motivations and Design Space Issues

In this section we discuss our design motivation and the design space issues in the core HyFinBall interface, i.e. the 6DOF and planar-3DOF modes and auto-mode switching.

3.3.1 Fatigue – Elbows-Resting vs Hands-Resting

Shaw and Green [8] advocate adding a user adjusted translational offset between the 6DOF button device and the 3D cursor in their two-handed system. This allows the user to keep his elbows resting

in his lap, or on the desk or chair arm to combat the common fatigue problems in VR interfaces. We include this offset in all our systems. However, in our prior experimental work [6] and in our formative evaluation of the HyFinBall interface, we found that while keeping elbows resting on a surface reduces fatigue compared to the naïve ‘arm’s outstretched’ approach, this interface is still more fatiguing than using a mouse. With a mouse, the hand--not just the elbow--rests on a surface.

Rich data visualizations involve coordinated views of both 2D and 3D components. Therefore we developed the HyFinBall UI with auto-mode switching between 6DOF and planar-3DOF mode to allow the user to perform one (or two-handed) 3D interactions as well as 2D interactions with his hand(s) resting on the desk. When both button balls are in planar-3DOF mode, the UI is essentially a two-handed mouse interface [49] but with additional specializations for stereo discussed below.

In our in-progress user study, there is a 6DOF-mode-only condition (III) that uses image-plane interaction for 2D component manipulations. Based on prior experience, we hypothesize that this condition will be more fatiguing than the HyFinBall interface (I) when the user must perform both 2D and 3D tasks or just 2D tasks. Of course, there is a trade-off. In the HyFinBall interface the user *must* place his hand on the table for 2D interaction. Pending our experimental results, it is likely important to test an approach that includes the HyFinBall auto-mode switch but also allows image-plane interaction for 2D interaction. An interesting question if a UI supports both options with what frequency does a user use each option?

The overall effect of these design space options, such as conditions I through IV, on fatigue and speed of user interaction will undoubtedly depend on the balance between the 2D interactions and 3D interactions used in a given application and the *temporal sequencing* and *period lengths* of planar-3DOF interactions and 6DOF interactions.

3.3.2 Auto-Switching between 2D and 3D cursors and 2D and 3D ITs

Display of a desktop GUI mouse cursor on top of a stereo image with positive and negative parallax (or even positive parallax alone) is well-known to aggravate stereo fusion problems. The HyFinBall interface addresses stereo cursor display issues in several.

In the HyFinBall interface, 3D cursors are displayed during 6DOF mode interactions while 2D cursors are displayed during planar-3DOF mode. We assume that during 2D interaction the user’s eyes fixate on geometry with zero parallax (i.e. the 2D graphical components) and that the user is not attempting to fixate on geometry with non-zero parallax. (The latter is the condition under which the naïve display of desktop 2D cursor creates problems).

Our anecdotal experience indicates this is the case, but future experimentation using an eye tracker could confirm this. Note, that we nonetheless, render the 2D cursors as slightly transparent discs so the user can see through them to any farther 3D geometry. Design space questions include the shape and transparency of the 2D cursors.

We have experimented with several additional visual cues when switching to planar-3DOF mode. First, we have experimented with enabling a simulation of depth-of-field image blur of the 3D geometry during planar-3DOF 2D interactions. The design space includes the presence/absence of the enabling of depth-of-field simulation and the fidelity of the depth-of-field rendering and its reduction of frame-rate. Second, we have experimented with reducing the eye separation when in planar-3DOF mode. If one HyFinBall is in planar-3DOF mode and is performing 2D interaction, then the modelled eye separation is cut in half. If both HyFinBall’s are in planar-3DOF mode and performing 2D interactions, eye separation is set to zero. The eye separation changes are animated over a 2s time period recommended by Ware et al [50].

Design space issues include presence/absence of the eye separation adjustment, the degree of adjustment, the rate of

adjustment, the conditions of adjustment and interaction with depth-of-field implementation. In general, our anecdotal results indicate eye separation reduction is useful when the user is performing a planar-3DOF 2D interaction.

Some planar-3DOF ITs are not strictly 2D interactions. In this case, the HyFinBall does not display a zero-parallax 2D cursor, but instead displays a “projected 3D cursor”. In a minor mouse version [51] of our prior travel technique for global terrain on a virtual workbench [48], we replaced the display of the desktop 2D mouse cursor with 3D cursor whose position was the intersection of a raycast from the cyclopean eye through the GUI’s reported mouse location. The projected 3D cursor can appear at any screen parallax depending on the location of the intersected terrain under the GUI cursor position. This approach is sometimes referred as geometry-sliding [52].

In DIEM-VR, we borrow this technique for the 2D lasso terrain selection. We chose for the planar-3DOF mode to perform the lasso operation rather than using a 6DOF mode image-plane technique. In the 2D lasso selection, the UI displays a projected 3D cursor (a cone) at position where a cast ray intersects the terrain. During 2D lasso selection we assume the user is fixating on the terrain surface location under the 3D cursor so the eye separation is set to its default setting. Our anecdotal experience strongly indicates this assumed fixation point is correct. An experimental evaluation with an eye tracker could confirm this.

These above examples indicate that there are important aspects to be considered when using the planar-3DOF mode on stereoscopic systems. There has been a fair amount of prior work in desktop 2D GUI’s regarding having the 2D image of the cursors change to indicate different application states or interaction modes. There has been interesting work in cursors for 3D selection such as Ware and Lowther’s One-Eyed cursor [53]. Teather and Stuerzlinger compared 4 cursor selection techniques in a [52], and more recently Bruder et al [54] explore different offset techniques on a virtual workbench. The HyFinBall raises additional questions because the cursor automatically switches between a 6DOF 3D cursor, a 2D zero-parallax cursor, and a projected 3D cursor (as in HyFinBall 2D lasso mode).

3.3.3 Form Factor

The third author and colleagues used dual, small form-factor (4.5 cm) button balls for multiple experiments [6] [57] in volume selection. In these experiments, we observe the typical adult participant switched between a more static grasp and a looser grasp. In the more static grasp, the user holds the 4.5 cm ball with the pinky, ring finger, middle finger against the thumb. The thumb is used for button presses (Figure 6A). In the looser grasp, all 5 fingers can roll the ball through all 3 rotation DOFs with the restriction of cord entanglement. We refer to the loose grasp as the rolling grasp, although it is understood that the hallmark of this “grasp” is its dynamic nature. The capability for the rolling grasp motion is what motivated Zhai’s to first use this size in his FingerBall. Regarding the more static grasp, Shaw and Green [8] discuss in detail the different grasps that can be adopted when 3 buttons are directly mounted on a Polhemus receiver.

Our observation of participants in experiments [6] [57] and our more recent anecdotal experience with our button balls in the HyFinBall UI, indicate that the particulars of the activated 3D IT dictates whether the static grasp or rolling grasp is best and that with minimal training a user can naturally switch between these two grasps when activating different 3D ITs.

For example, in the HyFinBall UI holding ButtonA activates a translation only scene-in-hand metaphor while holding ButtonC activates a 3D cursor centered view scaling [47]. Pressing ButtonB locks and activates a rotation-only scene-in-hand metaphor. Pressing any other button de-activates this mode. The first two IT’s work well with a static grasp. The rotation-only metaphor can be used with a static grasp, but performs better with the rolling grasp.

Switching to planar-3DOF mode requires placing the button ball on the desk with the buttons facing upward for thumb and pointer finger reach (Figure 6B). At present DIEM-VR does not use the yaw DOF.

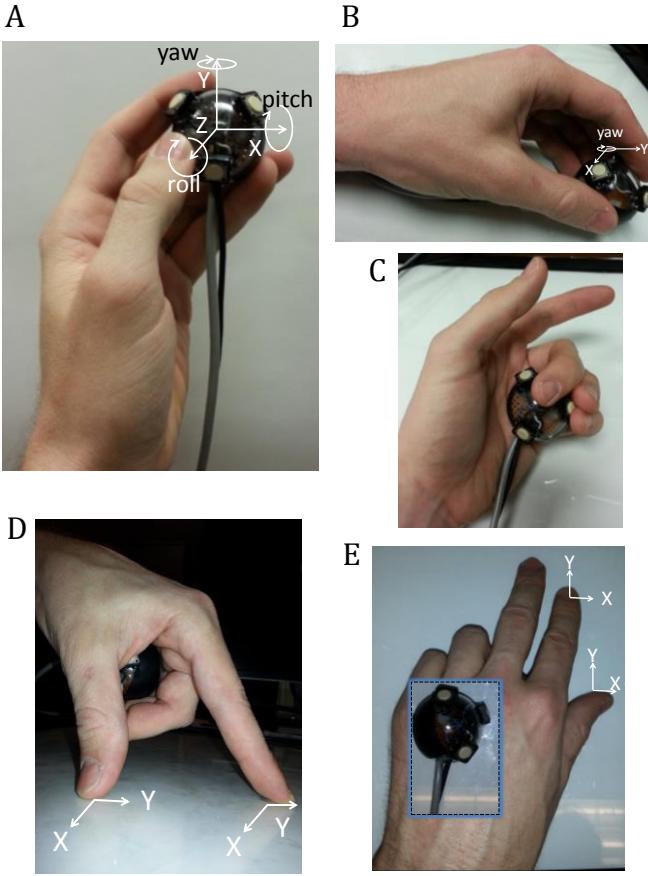


Figure 6A) Hand off table, 6DOF mode B) Hand on table, planar-3DOF mode. C) Dual fingers 3D gesture D) Fingers on table, multi-touch (side view). D) Fingers on table, multi-touch (top x-ray view showing held and hidden button ball).

3.4 HyFinBall+Finger-Tracking – The Interface

The HyFinBall+Finger-Tracking interface employs the previous described HyFinBall interface with 6DOF and planar-3DOF mode auto-switching while further leveraging the small-form factor and precision grip of the HyFinBall to allow the user's free fingers to interact with 3D finger tracking (Figure 1C and close-up Figure 6C) and 2D multi-touch (Figure 1D and close-up Figure 6D and E).

The touch detection is robust and we describe it's integration into DIEM-VR below. We have anecdotally found the 3Gear 3D finger tracking and gesture recognition less robust in its current iteration. In particular, we find a relatively high-number of false negatives during gesture recognition and find it's tracking range more limited than that of the Fastrak used for the HyFinBalls. However, the 3Gear system is relatively new and these issues vary with the variety of ways the Kinect's can be physically arranged.

The touch surface is a PQLab 24" 32-touch, multi-touch frame lying horizontally on the desk. A projector projects an image down onto the surface. A projector is needed instead of a flat panel display because when the button balls are used in planar-3DOF mode the display's metal ruins the EM tracking. Ideally a rear projected horizontal display would be used to avoid shadows, but in practice in this top-down configuration, the hands tends to cast projector

shadows in over nearly exactly the areas that are occluded from the users viewpoint.

DIEM-VR displays the Boolean query tree (Figure 4B) on the horizontal multi-touch display. DIEM-VR currently uses the horizontal display for pure-2D interactions, whereas the vertical display is used for previously described mix of 2D and 3D interactions. We specifically chose to touch enable the horizontal display rather than the vertical one, to maintain a hands-resting posture during the multi-touch interaction. Again, a key aspect of the HyFinBall UI design is that the user can use this thumb and pointer finger of one or both hands without dropping the HyFinBalls. Our anecdotal observation is that after 15-20 minutes of training, this can lead to a very fluid combined 2D and 3D interactions that would not be possible if user had to put-down and reacquire the HyFinBall. The query tree UI allows saving, restoring and deletion of queries and modification of query operators. Changes are immediately reflected in the terrain vertex highlighting, line net display and scatter plot highlighting.

3.5 HyFinBall+Finger-Tracking – Design Motivations and Design Space Issues

We use a horizontal, multi-touch display for the pure-2D interactions, while the planar-3DOF technique allows manipulation of 2D GUI components on the vertical display while keeping the hands resting.

The PQLab's frame generates touches due a HyFinBall placed on the table. The UI disables multi-touch when a HyFinBall is in planar-3DOF mode. Hence, multi-touch interaction must be done with the palm and HyFinBall off the table such as in Figure 6D, but this is fairly common on touch tables. Further, since we register the PQLabs touch surface with Fastrak tracked space it is theoretically possible to enable multi-touch when the HyFinBall is resting on the table by discounting the multi-touch touches that correspond to the HyFinBall position on the desk.

We integrate the 3Gear 3D hand and finger tracking space with the tracked workspace of the Polhemus tracker and PQLab's surface. The 3D hand tracking is robust enough to allow the UI to distinguish which hand generates a particular touch as long as the fingers of the two hands do not get too close together. Presently, we do not find the 3D finger tracking robust enough to add finger identification to the PQLab's touches nor to act as a substitute for the PQLab's multi-touch frame.

Anecdotally, the 3Gear system only tracks in a smaller working volume than that which one naturally uses with the Fastrak-based HyFinBall. We have experimented with detecting a pointing gesture and displaying a red ray emanating from either index finger. If made more accurate and wider in tracking range, this could provide ray based interaction that is enabled by finger gesture alone, rather than adding a ray-based interaction mode to the HyFinBall's. However, at present the loss of tracking and the error rate in gesture recognition preclude performing a usability study. Of course, this may improve with newer software releases.

A system with robust 3D finger tracking would make for an interesting design space of what interactions are best performed with hand+finger tracking and what are best performed with the HyFinBall's. Moehring and Froehlich performed a study using very robust and accurate hand and finger tracking by Vicon marked gloves compared with a 6DOF held-device (a Flystick) for a series of 3D manipulation tasks. Users preferred the naturalness of finger tracking. However, users of the Flystick performed significantly faster than "bare" finger tracking. Adding pinch-sensitive finger tracking improved task performance times to be within 10-20% of the Flystick condition.

We conjecture that for visualization applications that are used everyday, a minimally immersive and minimally encumbering system better fits many domain's users' desktop workflow. Markerless, head-tracked, auto-stereo displays has been demonstrated [58]. Non-encumbering input further implies using held devices and/or using gloveless hand and finger tracking which

precludes haptic feedback. Without any haptic feedback, it seems likely that held-devices will outperform 3D hand/finger tracking in 6DOF docking tasks while 3D hand/finger tracking could still enable other useful forms of 3D interaction. Hence, we suggest there is an interesting design space in hybrid 3D interfaces that support both held-device and hand/finger tracked 3D interaction.

4 FUTURE WORK

We are pilot testing a study that compares the following:

- I. the auto-switching HyFinBall UI
- II. dual planar-3DOF mode only UI
- III. dual 6DOF mode only UI
- IV. a single mouse

A participant performs a variety of 2D tasks, 3D tasks, and combination 2D followed by 3D tasks using DIEM-VR. All conditions use the desktop VR environment. The mouse mode uses projected 3D cursors for all 3D interactions and is our base-line condition. Our hypothesis is that after a training period, the core HyFinBall interface will perform better (faster) at the combined 2D+3D tasks than all other UIs and equal to the planar-3DOF only mode and mouse for 2D only tasks and equal to the 6DOF mode for 3D only tasks. We expect the 6DOF-only mode to be generally worse for 2D tasks and more fatiguing.

HyFinBall + multi-touch works robustly, but we have yet to formulate the user studies. HyFinBall + 3D hand/finger tracking is not yet robust enough to formally evaluate. Future multi-Kinect configurations may solve the problems with error rate and limited tracking range. Alternatively robust marker based hand and finger tracking could be employed. We are currently adding a dual camera Vicon system to our hardware ensemble. We believe there is an interesting design space to be explored when combining the HyFinBall UI with robust finger tracking.

Based on the ALCOVE [55] and our work in [49], we are in the process of configuring our system into a more seamless L-shaped display that also displays stereo 3D on the horizontal surface. Several of the prior works mentioned in Section 2 have begun exploring stereo + multi-touch, but to our knowledge prior work is limited.

ACKNOWLEDGMENTS

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

WORKS CITED

- [1] D. A. Bowman, E. Kruijff, J. J. LaViola and I. Poupyrev, 3D User Interfaces: Theory and Practice, Addison Wesley, 2005.
- [2] D. P. Mapes and J. M. Moshell, "A Two Handed Interface for Object Manipulation in Virtual Environments," *Presence*, vol. 4, no. 4, pp. 403-416, 1995.
- [3] C. Ware and D. Jessome, "Using the bat: a six-dimensional mouse for object placement," *Computer Graphics and Applications, IEEE*, vol. 8, no. 6, pp. 65-70, 1988.
- [4] B. Froehlich, J. Plate, J. Wind, G. Wesche and M. Gobel, "Cubic-Mouse-based interaction in virtual environments," *Computer Graphics and Applications, IEEE*, vol. 20, no. 4, pp. 12-15, 2000.
- [5] M. Deering, "High resolution virtual reality," in *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, New York, NY, USA, 1992.
- [6] A. Ulinski, C. Zanbaka, Z. Wartell, P. Goolkasian and L. F. Hodges, "Two handed selection techniques for volumetric data," in *3D User Interfaces, 2007. 3DUI'07. IEEE Symposium on*, 2007.
- [7] S. Zhai, P. Milgram and W. Buxton, "The influence of muscle groups on performance of multiple degree-of-freedom input," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 1996.
- [8] C. Shaw and M. Green, "Two-handed polygonal surface design," in *Proceedings of the 7th annual ACM symposium on User interface software and technology*, New York, NY, USA, 1994.
- [9] K. Hinckley, M. Sinclair, E. Hanson, R. Szeliski and M. Conway, "The VideoMouse: a camera-based multi-degree-of-freedom input device," in *Proceedings of the 12th annual ACM symposium on User interface software and technology*, New York, NY, USA, 1999.
- [10] L. Inc., "2D/6D Mouse technical reference manual," 1991.
- [11] D. Bowman, S. Coquillart, B. Froehlich, M. Hirose, Y. Kitamura, K. Kiyokawa and W. Stuerzlinger, "3D User Interfaces: New Directions and Perspectives," *Computer Graphics and Applications, IEEE*, vol. 28, no. 6, pp. 20-36, 2008.
- [12] J. Liu, S. Pastoor, K. Seifert and J. Hurtienne, "Three-dimensional PC: toward novel forms of human-computer interaction," in *In Three-Dimensional Video and Display Devices and Systems SPIE*, 2000.
- [13] K. Hinckley, "Input technologies and techniques," *The human-computer interaction handbook: fundamentals, evolving technologies and emerging applications*, pp. 151-168, 2002.
- [14] E. Kruijff, "Unconventional 3D User Interfaces for Virtual Environments," 2006.
- [15] T. Grossman, R. Balakrishnan, G. Kurtenbach, G. Fitzmaurice, A. Khan and B. Buxton, "Interaction techniques for 3D modeling on large displays," in *Proceedings of the 2001 symposium on Interactive 3D graphics*, New York, NY, USA, 2001.
- [16] W. Buxton and B. Myers, "A study in two-handed input," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 1986.
- [17] R. Balakrishnan and G. Kurtenbach, "Exploring bimanual camera control and object manipulation in 3D graphics interfaces," in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, New York, NY, USA, 1999.
- [18] R. Balakrishnan and K. Hinckley, "Symmetric bimanual interaction," in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, New York, NY, USA, 2000.
- [19] Y. Guiard, *Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model*, vol. 19, 1987, pp. 486-517.
- [20] W. Buxton, "Lexical and pragmatic considerations of input structures," *SIGGRAPH Comput. Graph.*, vol. 17, no. 1, pp. 31-37, 1983.
- [21] J. D. Mackinlay, S. K. Card and G. G. Robertson, "A Semantic Analysis of the Design Space of Input Devices," *Human-Computer Interaction*, vol. 5, pp. 145-190, 1990.
- [22] I. S. MacKenzie, R. W. Soukoreff and C. Pal, "A two-ball mouse affords three degrees of freedom," in *CHI '97 Extended Abstracts on Human Factors in Computing Systems*, New York, NY, USA, 1997.
- [23] M. W. Krueger, T. Gionfriddo and K. Hinrichsen, "VIDEOPLACE-an artificial reality," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 1985.
- [24] R. Wang, S. Paris and J. Popović, "6D hands: markerless hand-tracking for computer aided design," 2011.
- [25] K. Hinckley, R. Pausch, J. C. Goble and N. F. Kassell, "Passive real-world interface props for neurosurgical visualization," 1994.
- [26] M. Moehring and B. Froehlich, "Effective manipulation of virtual objects within arm's reach," March.
- [27] J. Mulder and R. v. Liere, "Personal Space Station," 2002.
- [28] H. Ishii and B. Ullmer, "Tangible bits: towards seamless interfaces between people, bits and atoms," 1997.
- [29] X. Bi, T. Grossman, J. Matejka and G. Fitzmaurice, "Magic desk: bringing multi-touch surfaces into desktop work," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2011.
- [30] G. Welch and E. Foxlin, "Motion tracking: no silver bullet, but a respectable arsenal," *Computer Graphics and Applications, IEEE*, vol. 22, no. 6, pp. 24-38, 2002.
- [31] H. Benko, E. Ishak and S. Feiner, "Cross-dimensional gestural interaction techniques for hybrid immersive environments," 2005.

- [32] M. Massink, D. Duke, S. Smith and V. S. Maria, *Towards Hybrid Interface Specification for Virtual Environments*, 1999.
- [33] M. F. C. Alencar, A. B. Raposo and S. D. J. Barbosa, "Composition of HCI evaluation methods for hybrid virtual environments," in *Proceedings of the 2011 ACM Symposium on Applied Computing*, New York, NY, USA, 2011.
- [34] F. Althoff, G. McGlaun, B. Schuller, P. Morguet and M. Lang, "Using multimodal interaction to navigate in arbitrary virtual VRML worlds," in *Proceedings of the 2001 workshop on Perceptive user interfaces*, New York, NY, USA, 2001.
- [35] D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiropoulos and D. F. Keefe, "Slice WIM: a multi-surface, multi-touch interface for overview+detail exploration of volume datasets in virtual reality," in *Symposium on Interactive 3D Graphics and Games*, New York, NY, USA, 2011.
- [36] P. Cohen, D. McGee, S. Oviatt, L. Wu, J. Clow, R. King, S. Julier and L. Rosenblum, "Multimodal interaction for 2D and 3D environments [virtual reality]," *Computer Graphics and Applications, IEEE*, vol. 19, no. 4, pp. 10-13, 1999.
- [37] H. Holzapfel, K. Nickel and R. Stiefelhagen, "Implementation and evaluation of a constraint-based multimodal fusion system for speech and 3D pointing gestures," in *Proceedings of the 6th international conference on Multimodal interfaces*, New York, NY, USA, 2004.
- [38] R. A. Bolt, "'Put-that-there': Voice and gesture at the graphics interface," in *Proceedings of the 7th annual conference on Computer graphics and interactive techniques*, New York, NY, USA, 1980.
- [39] P. Agrawal, I. Rauschert, K. Inochanon, L. Bolelli, S. Fuhrmann, I. Brewer, G. Cai, A. MacEachren and R. Sharma, "Multimodal interface platform for geographical information systems (GeoMIP) in crisis management," in *Proceedings of the 6th international conference on Multimodal interfaces*, New York, NY, USA, 2004.
- [40] I. Rauschert, P. Agrawal, R. Sharma, S. Fuhrmann, I. Brewer and A. MacEachren, "Designing a human-centered, multimodal GIS interface to support emergency management," in *Proceedings of the 10th ACM international symposium on Advances in geographic information systems*, New York, NY, USA, 2002.
- [41] J. Bouchet and L. Nigay, "ICARE: a component-based approach for the design and development of multimodal interfaces," in *CHI '04 Extended Abstracts on Human Factors in Computing Systems*, New York, NY, USA, 2004.
- [42] G. Bianchi, C. Jung, B. Knoerlein, G. Szekely and M. Harders, "High-fidelity visuo-haptic interaction with virtual objects in multi-modal AR systems," in *Mixed and Augmented Reality, 2006. ISMAR 2006. IEEE/ACM International Symposium on*, 2006.
- [43] R. Azuma, Y. Baillot, R. Behringer, S. Feiner, S. Julier and B. MacIntyre, "Recent advances in augmented reality," *Computer Graphics and Applications, IEEE*, vol. 21, no. 6, pp. 34-47, 2001.
- [44] D. Schmalstieg, A. Fuhrmann, G. Hesina, Z. Szalavári, L. M. Encarnaçāo, M. Gervautz and W. Purgathofer, "The studierstube augmented reality project," *Presence: Teleoper. Virtual Environ.*, vol. 11, no. 1, pp. 33-54, #feb# 2002.
- [45] C. Pirchheim, D. Schmalstieg and A. Bornik, "Visual Programming for Hybrid User Interfaces," in *Proc 2nd International Workshop on Mixed Reality User Interfaces*, 2007.
- [46] P. Song, W. B. Goh, C.-W. Fu, Q. Meng and P.-A. Heng, "WYSIWYF: exploring and annotating volume data with a tangible handheld device," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2011.
- [47] T. Butkiewicz, R. Chang, W. Ribarsky and Z. Wartell, "Understanding Dynamics of Geographic Domains," K. S. Hornsby, Ed., CRC Press/Taylor and Francis, 2007.
- [48] A. Ulinski, Z. Wartell, P. Goolkasian, E. Suma and L. Hodges, "Selection performance based on classes of bimanual actions," in *3D User Interfaces, 2009. 3DUI 2009. IEEE Symposium on*, 2009.
- [49] E. Houtgast, O. Pfeiffer, Z. Wartell, W. Ribarsky and F. Postl, "Navigation and Interaction in a Multi-Scale Stereoscopic Environment," 2005.
- [50] W. Robinett and R. Holloway, "Implementation of flying, scaling and grabbing in virtual worlds," 1992.
- [51] Z. Wartell, L. Hodges and W. Ribarsky, "Third-Person Navigation of Whole-Planet Terrain in a Head-Tracked Stereoscopic Environment," 1999.
- [52] R. C. Zeleznik, A. S. Forsberg and P. S. Strauss, "Two pointer input for 3D interaction," 1997.
- [53] C. Ware, C. Gobrecht and M. A. Paton, "Dynamic Adjustment of Stereo Display Parameters," *IEEE Transactions on Systems, Man and Cybernetics Part A: Systems and Humans*, vol. 28, no. 1, pp. 56-65, 1998.
- [54] Z. Wartell, "Stereoscopic Head-Tracked Displays: Analysis and Development of Display Algorithms," 2001.
- [55] R. Teather and W. Stuerzlinger, "Pointing at 3D targets in a stereo head-tracked virtual environment," 2011.
- [56] C. Ware and K. Lowther, "Selection using a one-eyed cursor in a fish tank VR environment," *ACM Trans. Comput.-Hum. Interact.*, vol. 4, no. 4, pp. 309-322, 1997.
- [57] G. Bruder, F. Steinicke and W. Sturzlinger, "Effects of Visual Conflicts on 3D Selection Task Performance in Stereoscopic Display Environments," 2013.
- [58] D. J. Sandin, T. Margolis, J. Ge, J. Girado, T. Peterka and T. A. DeFanti, "The VarrierTM autostereoscopic virtual reality display," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 894-903, 2005.
- [59] M. Meyer and A. Barr, "ALCOVE: design and implementation of an object-centric virtual environment," 1999.
- [60] L. D. Cutler, B. Frélich and P. Hanrahan, "Two-handed direct manipulation on the responsive workbench," 1997.