A Tangible Spherical Proxy for Object Manipulation in Augmented Reality

David Englmeier*
LMU Munich
Munich, Germany

Julia Dörner†
LMU Munich
Munich, Germany

Andreas Butz‡
LMU Munich
Munich, Germany

Tobias Höllerer§
University of California Santa Barbara
Santa Barbara, California

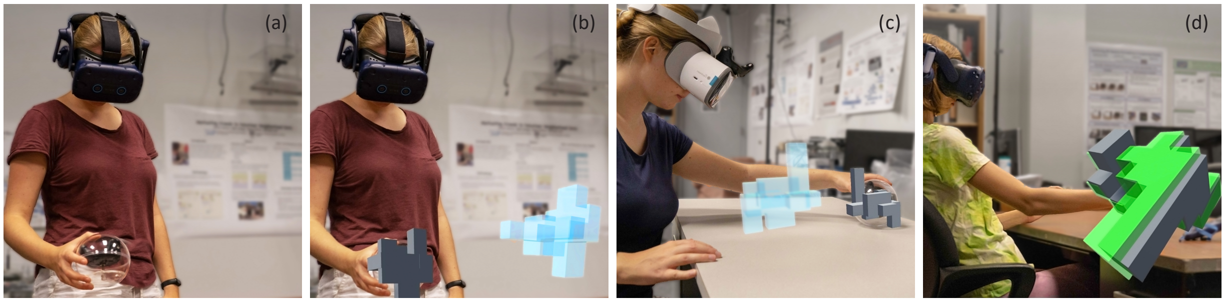

Figure 1: We investigate how a handheld sphere (a) aligned with an arbitrarily shaped virtual object (b) can serve as a universal tangible proxy for object manipulation in AR. We evaluated our concept by letting users perform an alignment task in mid-air (b) and on a table (c). To reach the given target (d), users had to perform the basic RTS (rotation, translation, scaling) operations.

Abstract

In this paper, we explore how a familiarly shaped object can serve as a physical proxy to manipulate virtual objects in Augmented Reality (AR) environments. Using the example of a tangible, handheld sphere, we demonstrate how irregularly shaped virtual objects can be selected, transformed, and released. After a brief description of the implementation of the tangible proxy, we present a buttonless interaction technique suited to the characteristics of the sphere. In a user study (N = 30), we compare our approach with three different controller-based methods that increasingly rely on physical buttons. As a use case, we focused on an alignment task that had to be completed in mid-air as well as on a flat surface. Results show that our concept has advantages over two of the controller-based methods regarding task completion time and user ratings. Our findings inform research on integrating tangible interaction into AR experiences.

Index Terms: Human-centered computing—Human computer interaction (HCI)—Interaction devices—Haptic devices; Human-centered computing—Human computer interaction (HCI)—Interaction paradigms— Mixed / augmented reality

1 Introduction

Augmented Reality (AR) can be seen as the real-time fusion of physical and virtual content in a 3D space [5], while Tangible User Interfaces (TUIs) allow the manipulation of virtual content by using physical objects [30]. One strength of TUIs is the close relation to already familiar properties of real-world items, such as physical characteristics and constraints, or a specific purpose of use defined by the object's shape [65]. This combination of physicality and naturalness can facilitate interfaces that are not only easy to use but also easy to learn and understand, due to their close resemblance to real-world interaction, and their literal direct manipulation [59].

As stated by MacIntyre [42], AR interfaces benefit from a strong relationship between the real and the virtual world. This is often achieved by mapping the input generated by physical objects to a virtual output [12]. Consequently, the concept of Tangible AR (TAR) has emerged [14, 35] as an obvious combination of both fields. Billinghurst et al. [13] define two properties of Tangible AR interfaces: the user interacting with virtual objects by manipulating a physical counterpart and each virtual representation being registered to one specific physical equivalent. One advantage of this approach is that an object can simply be selected by picking up the physical representation [37]. However, this can result in a significant demand for hardware, especially if interaction with a variety of different items is intended. A universal manipulator prop [38] used to select virtual objects can provide a solution to this problem. Selection thus becomes more scalable, but less intuitive. Also, the adaptation of physical to arbitrarily shaped virtual objects and rapidly changing virtual content is severely limited by physical constraints.

The development of tracking methods constitutes its own field in research, as they fundamentally enable AR interfaces. Most prominently, vision-based techniques, such as the tracking of markers or model-based tracking, are used to augment real items with virtual information and to track the 3D position, pose, or motion [37, 65]. However, if we interact naturally with real objects, our hands often cover large parts of the surface, rendering vision-based approaches difficult or impractical to implement, especially if fast and responsive tracking for resemblance to real physical behavior is desired.

AR and VR systems let users explore virtual environments naturally by controlling the camera with their head. High frame rates, responsive virtual scenes, and precise sensor-based tracking enable reactive AR applications with a large field of view. Our work uses video-see-through AR, utilizing standard VR headsets such as the tethered VIVE Pro and the wireless standalone Lenovo Mirage Solo. Both these systems support decent stereoscopic video feed-through modes. We see this technology as a good testbed for investigating natural interaction in AR due to its high-performance tracking capabilities and the possibility for opaque virtual augmentations.

An AR scene is likely to be populated with a variety of portable real-world objects. This raises the question of how to use such naturally existing objects in frequent interaction tasks such as 3D spatial manipulation. Our work investigates how simple, naturally shaped objects can support users during interaction in AR environments.

Spheres are found in many different contexts in our daily environments and constitute a fundamental shape that we encounter from an early age. We know how a sphere behaves, how it feels, and that we

*e-mail: david.englmeier@ifi.lmu.de
†e-mail: julia.doerner@campus.lmu.de
‡e-mail: butz@ifi.lmu.de
§e-mail: holl@cs.ucsb.edu

can rotate it easily no matter how we picked it up. If small enough, we can enclose it within our hands and fully perceive the shape even without looking at it. As stated above, the shape of a manipulator prop should ideally indicate its purpose or hint at how to use the object. Many natural objects can be picked up, rotated, and placed, but a sphere with its symmetrical shape and rotation may facilitate 3D object manipulation even better. Because most objects can not be scaled in the physical world, scaling operations are less obvious, but, as our work will show, can be modeled in accordance with the characteristics of the shape.

We propose a tool for 3D object manipulation, in which our proxy takes the role of an input device as well as a (simulated) display while its design is based on the related work now presented.

2 Background and Related Work

Our work builds on the fields of object manipulation and selection in AR and VR, tangible interaction on flat or interactive surfaces, and handheld spherical devices and displays.

2.1 Tangible Object Manipulation in VR and AR

In contrast to AR, where tracked physical counterparts of virtual objects often realize object manipulation, in Virtual Reality (VR), such tasks are likely to be performed with a dedicated controller that acts as a pointing device in 3D space, or by tracking the user's hands [15]. There are exceptions to this general pattern, such as *Handy AR* [39], which allows for hand interaction in AR, and several early VR systems that use physical props [28, 57].

An early example of using tracking gloves for 3D manipulation is the work by Mapes and Moshell [43]. They used two-handed interaction to realize basic tasks such as vertex manipulation, object transformation, or changing the viewpoint. Basic principles of this interface are still found in widespread manipulation techniques that are applied in many VR applications that rely on two-handed interaction, either with controllers or hand tracking. Another prominent example of two-handed object rotation using physical handles is the work of Ware et al. [64]. *Voodoo dolls* [49] is a two-handed VR interaction technique for manipulating objects at a distance. A shared physical device or interaction space can be used to collaboratively manipulate 3D content as described by Aguerreche et al. [1] or Duval et al. [21].

Comparing manipulation in AR and VR, Krichenbauer et al. [36] found that the AR condition led to lower task completion times for both mouse and a 3D input device. The AR application *iaTAR* [38] is based on tangible components, behaviors and intuitive interaction. A cube manipulator is used to select virtual objects by casting a ray. Once an object is selected, a distorted mapping is used for separating controls during interaction. This means that the virtual object stays in its position while the spatially separated manipulator defines the rotation. To deselect the object, the manipulator is hidden from the tracking device. Issartel et al. [31] extended the concept of a cubic AR device with a perspective corrected tangible AR display that allows for the inspection of contained content by rotation. The application *Tiles* [52] equally enables direct manipulation by strictly coupling one real cardboard marker to one virtual counterpart. Other examples use physical props such as a paddle or a scoop to pick up and translate virtual objects. [24, 35].

2.2 Object Selection and Scaling

Since our work is intended to resemble interaction with real objects while providing a reasonable demand in hardware expenses, selection techniques for grabbing virtual objects are highly relevant. In general we found two metaphors [4] for selection: virtual hand [51] and virtual pointing [40]. While the first relies on grabbing the target just as in the real world, the second allows for selecting objects out of the user's reach. Since the selection of small and distant objects is demanding by itself and not the focus of our work, we concentrate on classical virtual hand techniques based on direct manipulation. Mine et al. [44] describe several advantages related to a person's proprioception when operating within arms reach. When the object is within reach, selection is usually performed through (physical) buttons. However, dwell-time-based selection [61] may be a more general technique when using a natural object as a manipulator, because the object might lack physical properties that could serve as an input button. Automatic scaling is partially used to support selection techniques, for example, by increasing the size of pointing indicators or targets out of reach [3, 15, 50]. Scaling grabbed objects is often not supported in direct interaction techniques, because the six degrees of freedom (6-DOF) are logically assigned to rotation and translation, as demonstrated by the scene-in-hand technique by Ware and Osborne [63]. Two comparable metaphors to solve this problem were proposed by Cho et al. [19] with *Spindle+Wheel* and Song et al. [60] with the *Handlebar*. While the first utilized two spherical *Buttonball* devices [62] that simulated 7-DOF interaction, the latter relied on bi-manual hand gestures. Both techniques set the scale factor in dependence of the distance between the user's hands. Lastly, spheres also have a visual history in 3D manipulation tasks by serving as a widget for illustrating rotation [58].

2.3 Surface-Supported Object Manipulation

Augmented workbenches and interactive tables have a significant history of serving as a (shared) foundation for AR applications [8, 34, 35, 57]. Flat surfaces enable collaborative work by providing a shared interaction space. They allow the placement of physical props [16, 27, 47] and types of interaction that are not possible in mid-air, such as supported dragging or rolling. The *TableBall* [26] and the *Roly-Poly Mouse* [48] allow users to roll the input device. While the first is an example for manipulating content on an interactive surface at the device's current position, the second allows for a transition between 2D and 3D interaction in a desktop setup. It eventually outperformed established devices such as the Space Mouse by allowing for a selection of different configurations for the available six degrees of freedom. *TDome* [55] combined a flat multitouch surface with a base that could be rolled and subsequently fused physically with gesture-based input. Less frequently, non-planar surfaces are supplied to provide a base for AR applications, for example, in the form of virtual and AR sandtables [2, 54].

2.4 Handheld Spherical Devices

The most relevant category of related work for our concept includes fully spherical handheld devices that are (in contrast to several spherical displays or trackball devices) not mounted to a stand or a socket. Larger handheld spherical devices can be found in AR applications [10], either using outside projection or completely simulated in VR [22]. Louis and Berard demonstrated the superiority of an AR variant on a docking task that had to be completed by rotating the display [41]. Miyafuji et al. [45] realized a highly responsive [46] outside-projected spherical display, while the more portable, inside-projected approaches, have to rely on a socket to house projectors [7, 17]. Spherical displays that are small enough to fit into one hand have not yet been developed, but current VR technology provides an interesting opportunity for simulating them.

Finally, we extend our scope to devices that – to varying degrees – allow embracing a sphere, even if they can not display content on their surface. Early examples can be found in the work of Hinckley et al. [29] that describe benefits of handheld spherical devices in comparison to 2D input techniques while Poupyrev et al. [53] demonstrated the usability of non-isomorphic rotational mappings. Saidi et al. presented a palm-sized tangible device that was supported by rolling over the user's body [56]. The *GlobeFish* ported the concept of a 3D mouse to enable 6-DOF interaction with an almost entirely embraceable (potentially spherical) object that needed to be mounted to a suspension [25]. More distant examples for handheld devices

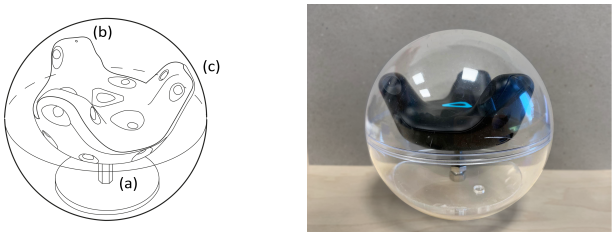

Figure 2: For a completely round shape, we screwed a circular plate (a) with a threaded rod to the bottom of a Vive Tracker (b) that was held in place by the top half of the sphere (c) pushing down on it. The right side of the figure shows the completely assembled device.

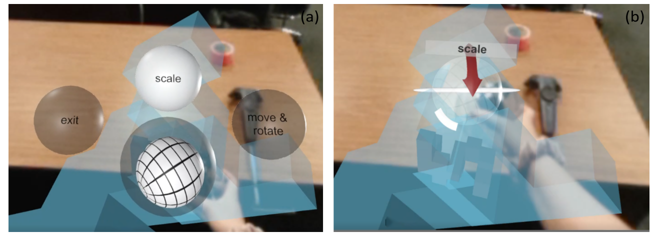

Figure 3: Since we did not intend the spherical device to have buttons, we used a menu (a) to switch between manipulation modes. We based selection and release on distance and dwell-time. In scaling mode (b), an arrow and an axis provided supporting visual cues.

that partially incorporated the characteristics of a sphere, particularly the tendency towards rotation, can be found in the work of Baudisch et al. [6] and Jackson et al. [32].

2.5 Summary

An interesting aspect of a spherical manipulator is the combination of a natural shape with tangible AR interaction and the intuitive interaction paradigm of literal direct manipulation. The resulting benefits may compensate for disadvantages [19], such as the lack of physical buttons for mode switching or simultaneous 7-DOF interaction. The wide applicability of the shape in 3D manipulation tasks, as well as the advantages provided by rolling interactions on an (augmented) surface, additionally support this assumption. Finally, we see potential in a handheld sphere acting as a collaborative device that can easily be shared among users by simply handing it over (without having to contemplate a certain orientation) and the interesting perspective that findings derived from an exploration of the device may transfer to other rotationally symmetric objects.

3 A Sphere for Tangible Interaction in AR

In this section, we will discuss the hardware construction and limitations of the spherical device as well as our design decisions for 3D spatial manipulation.

3.1 Device Construction

Our spherical device should be freely movable, lightweight, and provide the best possible tracking experience even if partially covered by both of the user's hands. In addition, it should provide a completely unobstructed surface to allow unhindered movement on a surface. We therefore decided to place a Vive Tracker[1] inside a two-piece acrylic glass sphere [23] as shown in Figure 2. We found that a 12 cm diameter would fit comfortably in one or two hands, and during bi-manual usage, the tracking experience would not noticeably decrease even if the sphere was embraced with both hands. The tracker is held in place by a socket that is fixated by the top half of the acrylic sphere pressing down on it. Consequently, the tracker does not sit exactly in the center, but the outside surface of the sphere remained completely untouched. Due to its overall light weight (ca. 190g), the resulting slight imbalance was hardly noticeable even when the device was fiercely rotated. The construction of the device is cheap and simple, allowing the production of multiple devices that could be coupled via positional and gestural relationships [13].

3.2 Interaction Design

We based our main design decisions on two requirements. First, interaction with the sphere should be possible without any physical or simulated buttons, only by means of direct manipulation. Second, the technique should take advantage of the spherical shape. These decisions reflect our focus on exploring the effects of the spherical shape, which would be shifted or complicated by adding buttons or a multi-touch surface. The exploration of such interaction possibilities might be interesting in its own right, but currently is made difficult by hardware limitations, especially for a spherical surface [22]. Our approach has the advantage of keeping expenses at a minimum while maintaining a simple and elegant object. The work highlights the usability and acceptance possibilities of buttonless interaction devices. Under these constraints we decided to implement two different transformation modes: *Translate + Rotate* (6-DOF) and *Scale* (1-DOF). To switch between the modes, we used a simple menu (Figure 3, (a)) that appears when the user selects a virtual object. We realized selection and release of an object based on distance and dwell-time. For scaling, we also experimented with a distance-based approach that we compared to a rotation-based approach (Figure 3, (b)). Due to the self-imposed constraint to use no buttons and the resulting impossibility for "clutching", as well as the round shape, we decided in favor of the latter. We will report further details on the implementation of the interaction modalities in Section 4.3.1. In line with manipulation techniques leveraging physical manipulators, we decided to couple the virtual and physical objects at their center of gravity or respective pivot point [15, 38].

4 Experiment

To evaluate our design, we compared it in a lab study to three manipulation techniques using the HTC Vive Controller[2], including state-of-the-art object manipulation mapping using the controller hardware trigger button and trackpad.

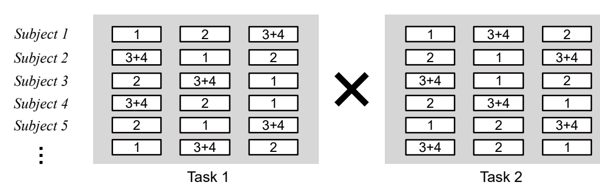

Figure 4: To reduce adverse order effects, we counterbalanced the sequence of conditions between subjects as well as between tasks. Note that Condition 4 was always presented after Condition 3. We randomly selected 30 out of 36 possible orderings.

4.1 Study Design

We recruited 30 participants (15 male, 15 female) with an age of 18 to 32 years (mean: 22.3, SD: 4.26). Subjects had a mean self-reported expertise with VR of 3.2 (from 1 = no experience, to 10 = expert). The study was executed as a within-subjects experiment with the manipulation technique (1: *Sphere*, 2: *Controller: Buttonless*, 3: *Controller: Button and Menu*, and 4: *Controller: Button and*

[1]https://www.vive.com/us/vive-tracker/

[2]https://www.vive.com/us/accessory/controller/

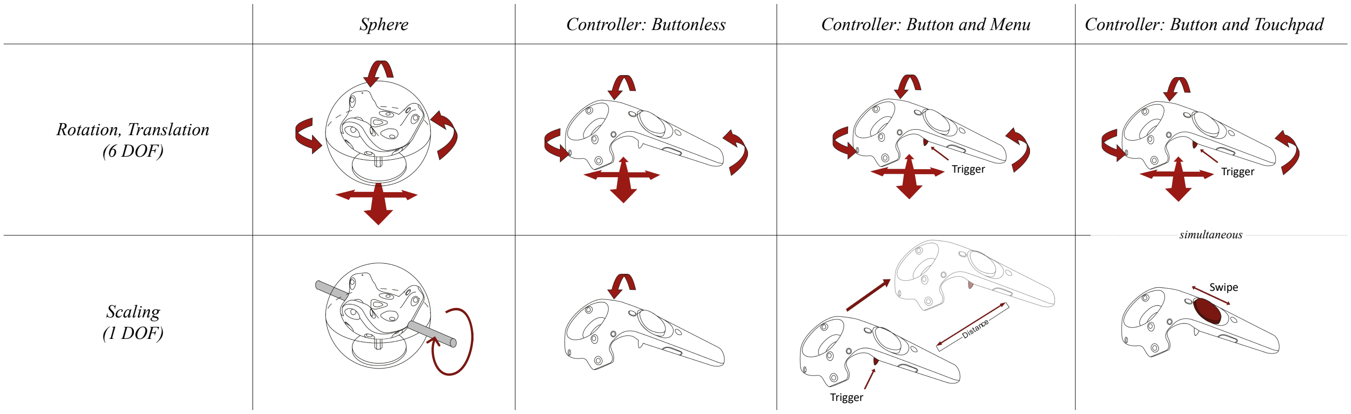

Figure 5: For the Sphere condition, the prop's six degrees of freedom are directly mapped to the virtual object. Scaling is performed by rotating the sphere around one axis. The second condition (controller-as-sphere) works in the same way but with a different physical device (VIVE controller). The third condition (controller-with-trigger) makes use of the controller's trigger button, which enables "clutching" and one-handed rotation. The last condition (controller-trigger-touchpad) allows for simultaneous 7-DOF manipulation by employing a trigger button and trackpad.

Touchpad) as the independent variable and two task scenarios: mid-air manipulation and table-top manipulation. The order of conditions was permuted using a counterbalanced design as illustrated in Figure 4. We defined the constraint that the fourth condition (*Controller: Button and Touchpad*) that we considered as an evolution of the third (*Controller: Button and Menu*) would always be completed en bloc and also permuted the succession of conditions between both tasks. This resulted in a total of $6 \times 6 = 36$ permutations. Each of the scenarios had a sequence of object alignment challenges, each manipulation technique was given carefully administered scripted training time before it was tested, and the whole study took about 50 minutes (including a post-experiment questionnaire). The subjects were compensated for their time with a $10 payment.

4.2 Apparatus

Apart from the already described tracked spherical device, we used additional hardware and software to implement our design and to realize the experimental setup.

4.2.1 Hardware

In terms of HWDs (Head-Worn Displays) we tested two devices: the *Lenovo Mirage Solo with Daydream* (Figure 1, (c)) and the *HTC Vive Pro Eye* as illustrated in Figure 1, (a, b, d). The first is a standalone headset that does not require external tracking equipment and provides a 75 Hz refresh rate and a 110° field of view (FOV) while the second requires a permanent (wired) connection to a PC and provides a refresh rate of 90 Hz and an equally sized FOV. We ultimately decided to use the HTC Vive due to the better quality of the colored stereoscopic camera image and the easier integration of the object tracking system, that required an external tracking server for the Lenovo HWD.

4.2.2 Software

To realize the virtual environment for our study, we used Unity[3] and C# as the programming language. To provide visual feedback for interaction with the sphere during rotation, we added a simple black and white grid and supplied a menu for interaction techniques that required mode switching (as seen in Figure 3, (a)). The menu would appear when approaching a virtual object with the device within a distance of less than 20 cm. We based the design of the menu on the concept of a 3D ring menu [20, 40] that we found to perform well with VR controllers and the sphere. It consisted of three spherical items (15 cm diameter) positioned in equal distance (20 cm) to the respective object's center. To select a menu item, it needed to be approached with the controller: the item closest to the input device was highlighted and if the distance was reduced to below 10 cm it was selected. We provide a detailed explanation of the manipulation modes and the implementation of object release along with the description of the four conditions in the next section.

4.3 Experimental Conditions

In order to detect the advantages and disadvantages of the spherical form factor, the dwell-time and time needed to switch modes, as well as the rotation-based scaling, we compared our technique to three controller-based methods that progressively increased hardware complexity (see Figure 5), resulting in the following study conditions:

4.3.1 Sphere

When the sphere is approaching the desired virtual object that the user intends to manipulate, the described mode menu appears. When selecting "move & rotate" the virtual representation of the sphere is changed to nearly transparent while the virtual object snaps to the center of the sphere, and subsequently, rotation and translation are directly mapped to the grabbed object. In "scale" mode, the rotation of the sphere around an indicated axis is mapped linearly to the scale of the virtual object. When users rotate the sphere towards themselves (counter-clockwise), the virtual objects increases in size. To exit the current manipulation mode, the sphere has to be kept still for the dwell-time of a second while the circular progress bar indicates the remaining time to complete the interaction. Now the menu reappears, and the user can select another mode or cancel interaction with selecting "exit".

4.3.2 Controller: Buttonless

The buttonless controller interface uses the same interaction technique as the sphere. The only difference is the controller's shape: a bare HTC Vive controller was employed instead of the sphere. None of its hardware controls were used, just its tracking capability.

4.3.3 Controller: Button and Menu

This interface adds the use of a button to the previous condition. It still utilizes the same menu, but we extended the manipulation methods by adding a physical button. In the "move & rotate" mode the trigger is used to grab and release the object. This allowed for one-handed use of the controller, especially during rotation. For scaling,

[3] https://unity.com/

we implemented a comparable "clutching" technique: Moving the controller away from the virtual object while the trigger is pressed increases the size of the virtual object while moving towards the object decreases the size. The functionality to completely release the object remained unchanged. This condition mirrors standard 'grab to move/rotate' and 'grab to scale' functionality from VR authoring tools, and represents best practices afforded by a tracked controller with a single button.

4.3.4 Controller: Button and Touchpad

As a mode-less interaction technique not suffering from the potential disadvantages of the previous ones, we implemented a technique using buttons and touchpad for providing simultaneous 7-DOF interaction. The object can be grabbed with the trigger as in the previous technique and scaled at the same time via the touchpad. Swiping up on the touchpad increases size while swiping down decreases the scale. Consequently, no menu or dwell-time is needed, a huge advantage afforded by the additional hardware controls. This condition mirrors best practices for the use of controller buttons and hardware for object manipulation from VR authoring tools.

4.4 Tasks

To test our conditions, subjects performed an alignment task, first in mid-air (task scenario 1), and then with the support of a table (task scenario 2). The design of the second scenario is motivated by possible effects the shape of the controller could generate in conjunction with a rigid surface that also provides ergonomic benefits and could potentially reduce fatigue effects occurring during the first task.

Before the actual tasks started, participants were shown a video of all conditions and had to complete a simpler, preliminary alignment task until they felt comfortable with each method. A short voice recording indicated the controller technique that had to be used when we applied a new condition. For each condition, a fixed succession of four objects (Figure 6) had to be aligned with respective targets. Each alignment was repeated once. To control for a learning effect, we permuted and balanced the sequence of objects while each object was associated with a predefined transformation (RTS).

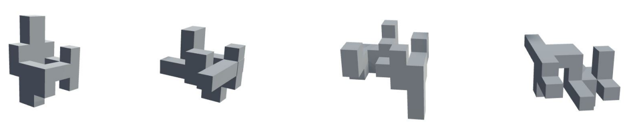

Figure 6: As primary objective, users had to align four different objects with a target. In a first task scenario, this was done in mid-air, in a subsequent second one, this was done on a tabletop.

During the task, we showed two versions of the same object and instructed the participant to place the gray version so that it fitted a blue transparent template. The target turned green to indicate that the placement/pose was sufficiently close, which occurred when the user reached predefined margins of error: for translation, the deviation needed to be smaller than 6 cm for rotation less than 10°, and in scale 30% above or below the target's size. After three seconds that could be used to enhance the result further, the next subtask was started and advised subjects to solve each task as quickly as possible, while, as a secondary objective aiming for an accurate placement within the defined time frame of three seconds. Hence, the process can be divided roughly into four phases: selection, inspection, manipulation, and placement.

We chose the margins of error and the task termination procedure to support an overall focus of the design on task completion time. These choices prevented users from spending a majority of the task time in the last phase performing final adjustments. Therefore, we expected users to concentrate on the main objective of solving the tasks quickly while the secondary objective of accurate placement was presented in a separate stage. These choices stemmed from our intended use cases in AR object manipulation, in which the main goal is to quickly grab, analyze, and place virtual objects via a tangible proxy. In related manipulation studies that focus more on the context of 2D desktop applications, self-defined termination is often applied [11, 18, 28].

Upon the completion of tasks, we asked participants to fill out a questionnaire. Considering seven different aspects, they had to rate each controller interface in the context of both tasks. Additionally, users had to choose their favorite interaction method and provide a reason why they chose this specific method. The last two questions allowed the participants to point out any aspect of a method that they liked or disliked.

4.5 Experimental Hypotheses

Regarding the four conditions, we formulated three hypotheses:

H_1: Users in the condition *Sphere* will perform significantly better than in the condition *Controller: Buttonless* due to rotational advantages of the sphere. If this hypothesis can be supported, it provides evidence for the spherical shape being superior to a rotationally asymmetric object such as the bare controller.

H_2: The condition *Controller: Buttonless* will provide inferior performance to the condition *Controller: Button and Menu*. If this hypothesis is affirmed we can deduce that the integration of a physical button and subsequent interaction techniques are beneficial even if the need to switch between modes still prevails. In relation to *Sphere*, we did not formulate a hypothesis regarding performance differences to *Controller: Button and Menu* since we believed that the respective advantages of either condition could make each one come out ahead.

H_3: The condition *Controller: Button and Touchpad*, which was designed as the full-hardware-support state-of-the-art comparison, will be superior to all other conditions. If this hypothesis is supported, it demonstrates the benefits of simultaneous 7-DOF interaction and the absence of mode switching.

5 Results

Below, we will present our quantitative data on task performance (completion time, accuracy) and the results from the post-experiment questionnaires followed by additional qualitative findings we recorded during our study.

5.1 Quantitative Results

As the main objective for recording user performance, we measured task completion time. The first task took subjects about 160-320 seconds, while the second took about 150-250 seconds. We additionally logged how precisely users managed to align the objects with targets taking into account the average of the normalized deviations in scale, rotation (angular difference) and translation. Finally, we present results from the post-experiment questionnaire.

5.1.1 Task Completion Time

To test for significant differences among the controller conditions, we ran a repeated measures ANOVA with multivariate evaluation. The Pillai's trace and Wilk's lambda test revealed statistical significance for the four input conditions in the first and second task: $F(3,27) = 22.85$, $p < 0.001$ and $F(3,27) = 23.68$, $p < 0.001$. We performed pairwise comparisons and used a Bonferroni-corrected t-test for post hoc analysis. We will, at first, focus on the results regarding the condition *Sphere*.

In the first task scenario (mid-air), the *Sphere* provided significantly lower task completion times than *Controller: Buttonless* and *Controller: Button and Menu*: $p = 0.038$ and $p < 0.001$. The condition *Controller: Button and Touchpad* outperformed the condition *Sphere* in this task with $p < 0.001$.

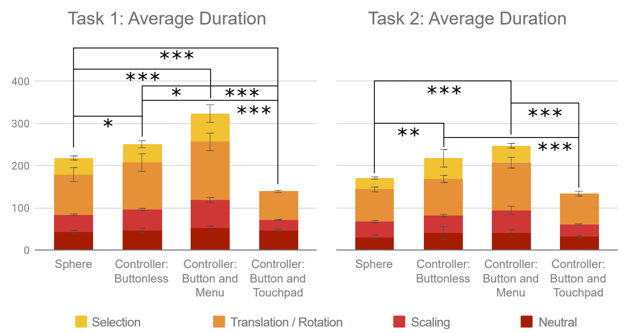

Figure 7: Total duration for both tasks. The four conditions are evaluated by accumulated average values of times spent for interaction types. Values are given in seconds with 95% confidence intervals.

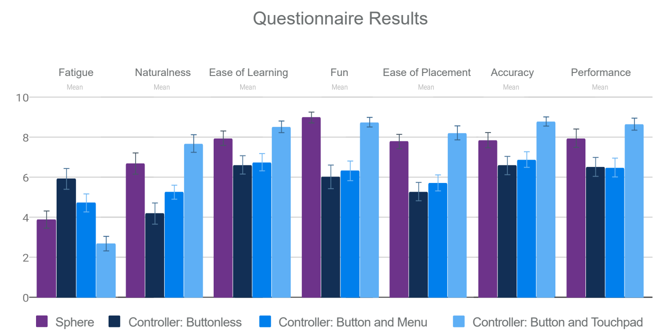

Figure 8: Users' perception ratings with 95% confidence intervals for both tasks. Ratings were given on a 10-point Likert scale. Apart from fatigue a higher rating represents a better result.

The results of the second task scenario (table-top) showed similar effects. Again, *Sphere* surpassed the conditions *Controller: Buttonless* and *Controller: Button and Menu* with: $p = 0.002$ and $p < 0.001$. However, the spherical controller in this task statistically did not differ significantly from the condition *Controller: Button and Touchpad* ($p > 0.999$).

For a deeper analysis regarding the *Sphere* condition, we compared the times that users spent in the different states during interaction. For the first task scenario, we found that the overall lower time for the condition *Sphere* in comparison to the second condition (*Controller: Buttonless*) was the result of significantly faster Scaling ($F(3,27) = 30.55$, $p < 0.001$). The third condition (*Controller: Button and Menu*) was outperformed in Scaling ($p < 0.001$), Translation/Rotation ($F(3,27) = 26.24$, $p < 0.001$) and Selection ($p = 0.007$). The lower overall completion times for the condition *Controller: Touchpad* in comparison to the *Sphere* were a result of significant effects in both Scaling ($p < 0.001$) and Translation/Rotation ($p = 0.002$).

In the second task scenario, *Sphere* surpassed *Controller: Buttonless* due to faster Scaling ($p = 0.303$) and Selection ($p = 0.021$) as well as *Controller: Button and Menu* because of quicker Selection ($p = 0.007$), Scaling ($p = 0.035$) and Translation/Rotation ($p < 0.001$).

A comparison among controller conditions revealed that *Controller: Touchpad* outperformed the other controller based conditions both with $p < 0.001$. Additionally, *Controller: Buttonless* conditioned significant lower task completion times than *Controller: Button and Menu*. Figure 7 provides an overview of all completion times for both task scenarios.

5.1.2 Accuracy

Although accuracy was not the main objective of the study, we completed a repeated measures ANOVA with multivariate evaluation on the mean deviations for aligning the virtual object with the target. As expected, the Pillai's trace and Wilk's lambda test revealed that the difference was not significant, neither for the first ($F(3,27) = 0.32$ $p = 0.81$) nor for the second task ($F(3,27) = 2.50$, $p = 0.081$).

5.1.3 Questionnaire Results

To evaluate the spherical device regarding subjective ratings, participants were asked to answer a post-experiment questionnaire (10-point Likert scale) for each task separately. We found that the ratings for both task scenarios were highly similar, hence we will present them in combination. Figure 8 provides an overview of both tasks. To reveal significant effects, we performed a Friedman test on the given ratings. We discovered a significant effect for all ratings and ran a Dunn-Bonferroni post hoc test for pairwise comparisons.

We found a significant influence of the conditions for question regarding fatigue ($\chi^2(3) = 29.692$, $p < 0.001$). The condition *Controller: Buttonless* was perceived as more fatiguing than *Sphere* ($p = 0.01$), and *Controller: Touchpad* ($p < 0.001$). *Controller: Touchpad* was additionally considered significantly less fatiguing than *Controller: Button and Menu* ($p = 0.01$). In terms of naturalness ($\chi^2(3) = 39.260$, $p < 0.001$) users rated the *Sphere* as significantly more natural than *Controller: Buttonless* ($p < 0.001$) and *Controller: Button and Menu* ($p = 0.01$) while *Controller: Touchpad* ($p < 0.001$) was perceived as more intuitive than the other controller conditions two and three ($p < 0.001$ and $p < 0.05$). Ratings for ease of learning ($\chi^2(3) = 24.466$, $p < 0.001$) put *Controller: Touchpad* in front of *Controller: Button and Menu* ($p = 0.01$) and *Controller: Buttonless* ($p < 0.001$) while the *Sphere* also surpassed the latter ($p < 0.05$).

In terms of perceived fun the *Sphere* and *Controller: Touchpad* were rated higher than *Controller: Button and Menu* and *Controller: Buttonless* with $p < 0.001$ and $p = 0.01$ for both comparisons. When asked which technique facilitated easy object placement ($\chi^2(3) = 46.236$, $p < 0.001$) the subjects found the *Sphere* to be superior to *Controller: Buttonless* ($p = 0.01$). *Controller: Touchpad* outperformed condition two and three both with $p < 0.001$.

Lastly, we asked about subjective performance and accuracy. For the latter ($\chi^2(3) = 32.891$, $p < 0.001$) participants felt to be more precise when using the *Sphere* rather than the buttonless controller ($p = 0.01$). Again, the controller with touchpad was rated significantly higher than both other controller conditions ($p < 0.001$). The perceived performance ($\chi^2(3) = 36.151$, $p < 0.001$) was rated significantly higher for the *Sphere* when compared to to condition two and three with $p = 0.01$. *Controller: Touchpad* achieved the same significance with $p < 0.001$.

5.2 Qualitative Results

During the experiment, we observed that most participants first tried to rotate and place the virtual object and then scale it. This resulted in longer completion times since it usually required re-adjusting. Users who started by scaling were faster for the most part. When operating with the sphere or the buttonless controller, users generally used both hands when performing rotations. For the condition, *Controller: Button and Menu* bi-manual interaction increased the difficulty to press the button and thus was rarely used.

We asked participants to choose their preferred controller from eight options, which described each controller condition for interaction in mid-air and on the table. 58.6% favored the condition *Controller: Button and Touchpad* in mid-air, 31% the *Sphere* in mid-air. Only one participant each voted for *Controller: Button and Menu* and *Controller: Buttonless*. Under the table constraint, the only condition to receive one vote was *Sphere*. Users who chose the

Sphere stated that it felt more natural and intuitive, especially for rotation. Participants opted for the *Controller: Button and Touchpad* because it felt more efficient and required no menu and dwell-time.

In the feedback section of the questionnaire, most users criticized the table, since it limited their movements and made the alignment task a lot harder. Additionally, they described *Controller: Button and Menu* as more complicated as the other interaction methods while *Controller: Buttonless* strained users' wrists, especially when operating it with only one hand.

6 Discussion and Outlook

Considering the above results, we can not only affirm H_1 but even extend it to the statement that the spherical device also could outperform the "clutching" techniques supported by the third condition. Its superior performance showed throughout the tasks interestingly most consistently influenced by quicker scaling. This leads to the insight that for a spherical object, scaling based on rotation is a viable approach, while we can attribute other positive effects in selection, translation, and rotation equally to the ergonomics of the device. The sphere's advantages regarding selection may also hint to an interesting benefit of the shape. The clearly recognizable center point of the object may have facilitated the selection of objects by approaching them. For the controller, the less clear center may have made the judgment of distances more difficult and could be an explanation for the inferior performance in selection. These findings are backed by the generally positive subjective ratings of the sphere that were on par with the most convenient hardware-supported fourth condition.

If we evaluate H_2, we have to reject the hypothesis partially. The increased hardware complexity for condition three (by adding a physical button) did, in case of task completion times, result in significant disadvantages. Additionally, we find in comparison to condition two similar ratings regarding accuracy, performance, fun, and ease of learning, indicating that the more sophisticated hardware does not positively affect these areas. However, users rated the third condition higher than the buttonless controller, which leads us to the assumption of this condition generating a high physical demand mainly due to the rotation of an asymmetrical object.

As expected, we can accept H_3. The simultaneous 7-DOF interaction and the time savings due to no mode switches led to results that could not be surpassed by any other condition. Yet, for the table-supported task, the only condition that was not significantly outperformed by *Controller: Button and Touchpad* was the sphere. The observation of the superiority of the 7-DOF condition should be seen in the context of directly coupling the virtual object to the center of the manipulator, while a changed perceptual space could lead to a different result [33].

The fact that users did not prefer working with the table was also expected. The surface, while possibly providing ergonomic benefits for long-term use, deliberately limited some pose choices and, therefore, occasionally complicated object placement. However, the rotationally symmetric sphere could not generate a significant advantage from this limitation. The situation that an object was scaled smaller than the physical sphere's diameter did not occur in our experiment. Still, this theoretical problem for placement on a surface appears to be solvable more easily in the case of the sphere. We implemented a simple function that would – depending on the distance of the sphere's center to the table – align the virtual object with the edge of the sphere that is closest to the surface, which enables the placement of small fully contained virtual objects.

While it was our goal to explore the advantages of a buttonless sphere, our findings indicate that a spherical manipulator may be desirable as a modeless interface that fully supports 7-DOF interaction. Although this goes beyond the scope of our work, we see potential in exploring the simulation of buttons or gesture-based interaction [9] on a handheld spherical proxy. Hardware extensions that could implement an additional degree of freedom are also conceivable.

7 Conclusion

We demonstrated that for the use case of a tangible AR manipulator with two separate modes, a handheld spherical device has significant advantages in comparison to an asymmetric controller in task completion time, in mid-air, on a surface, and in terms of user perception. Consistent with textbook knowledge, we saw that the modeless input condition, which allowed full 7-DOF operation simultaneously, outperformed all other conditions, but at the cost of the substantially more complex controller hardware. We, therefore, state that for an AR application situated in a natural environment, a tracked sphere provides a solid basis for 3D spatial manipulation.

In view of the high ratings in naturalness and ease of learning, AR environments that require fast comprehension of interaction with natural objects appear to offer particularly promising prospects.

Given our initial motivation of using different portable, familiarly shaped objects for controlling virtual object manipulations, there clearly remains work to be done, since this paper focused specifically on the benefits of spherical shapes as proxy controllers, and also did not explore the use of buttons or interaction surfaces on the sphere.

Yet, our findings clearly show benefits of this kind of symmetrical manipulation interface via physical proxy, which allows for bi-manual operation and implements scaling by rotation.

References

[1] L. Aguerreche, T. Duval, and A. Lécuyer. Reconfigurable tangible devices for 3d virtual object manipulation by single or multiple users. In *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*, pp. 227–230. ACM, 2010.

[2] C. R. Amburn, N. L. Vey, M. W. Boyce, and J. R. Mize. The augmented reality sandtable (ares). Technical report, 2015.

[3] F. Argelaguet and C. Andujar. Improving 3d selection in immersive environments through expanding targets. In *SG'08: Proceedings of the 9th international symposium on Smart Graphics*, pp. 45–57, 2008.

[4] F. Argelaguet and C. Andujar. A survey of 3d object selection techniques for virtual environments. *Computers & Graphics*, 37(3):121–136, 2013.

[5] R. T. Azuma. A survey of augmented reality. *Presence: Teleoperators & Virtual Environments*, 6(4):355–385, 1997.

[6] P. Baudisch, M. Sinclair, and A. Wilson. Soap: a pointing device that works in mid-air. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, pp. 43–46. ACM, 2006.

[7] O. Belloc, M. Nagamura, D. Fonseca, A. Rodrigues, D. Souza, C. S. Kurashima, M. Almeida, E. Z. Borba, R. Lopes, and M. K. Zuffo. Orbevr: a handheld convex spherical virtual reality display. In *ACM SIGGRAPH 2017 Emerging Technologies*, p. 19. ACM, 2017.

[8] H. Benko, R. Jota, and A. Wilson. Miragetable: freehand interaction on a projected augmented reality tabletop. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 199–208. ACM, 2012.

[9] H. Benko, A. D. Wilson, and R. Balakrishnan. Sphere: multi-touch interactions on a spherical display. In *Proceedings of the 21st annual ACM symposium on User interface software and technology*, pp. 77–86. ACM, 2008.

[10] F. Berard and T. Louis. The object inside: Assessing 3d examination with a spherical handheld perspective-corrected display. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 4396–4404. ACM, 2017.

[11] L. Besançon, P. Issartel, M. Ammi, and T. Isenberg. Mouse, tactile, and tangible input for 3d manipulation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 4727–4740, 2017.

[12] M. Billinghurst, A. Clark, G. Lee, et al. A survey of augmented reality. *Foundations and Trends® in Human–Computer Interaction*, 8(2-3):73–272, 2015.

[13] M. Billinghurst, R. Grasset, and J. Looser. Designing augmented reality interfaces. *ACM Siggraph Computer Graphics*, 39(1):17–22, 2005.

[14] M. Billinghurst, H. Kato, and I. Poupyrev. Tangible augmented reality. *ACM SIGGRAPH ASIA*, 7, 2008.

[15] D. A. Bowman and L. F. Hodges. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. *SI3D*, 97:35–38, 1997.

[16] A. Butz, O. Hilliges, L. Terrenghi, and D. Baur. Hybrid widgets on an interactive tabletop. *Ubicomp'07: Adjunct Proceedings, Demo session*, 2007.

[17] M. Cabral, F. Ferreira, O. Belloc, G. Miller, C. Kurashima, R. Lopes, I. Stavness, J. Anacleto, S. Fels, and M. Zuffo. Portable-spheree: A portable 3d perspective-corrected interactive spherical scalable display. In *2015 IEEE Virtual Reality (VR)*, pp. 157–158. IEEE, 2015.

[18] M. Chen, S. J. Mountford, and A. Sellen. A study in interactive 3-d rotation using 2-d control devices. In *Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pp. 121–129, 1988.

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

[20] R. Dachselt and A. Hübner. Three-dimensional menus: A survey and taxonomy. *Computers & Graphics*, 31(1):53–65, 2007.

[21] T. Duval, A. Lécuyer, and S. Thomas. Skewer: a 3d interaction technique for 2-user collaborative manipulation of objects in virtual environments. In *3D User Interfaces (3DUI'06)*, pp. 69–72. IEEE, 2006.

[22] D. Englmeier, I. Schönewald, A. Butz, and T. Höllerer. Feel the globe: Enhancing the perception of immersive spherical visualizations with tangible proxies. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1693–1698, March 2019.

[23] D. Englmeier, I. Schönewald, A. Butz, and T. Höllerer. Sphere in hand: Exploring tangible interaction with immersive spherical visualizations. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 912–913, March 2019.

[24] M. Fjeld and B. M. Voegtli. Augmented chemistry: An interactive educational workbench. In *Proceedings. International Symposium on Mixed and Augmented Reality*, pp. 259–321. IEEE, 2002.

[25] B. Froehlich, J. Hochstrate, V. Skuk, and A. Huckauf. The globefish and the globemouse: two new six degree of freedom input devices for graphics applications. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pp. 191–199. ACM, 2006.

[26] M. Hancock, O. Hilliges, C. Collins, D. Baur, and S. Carpendale. Exploring tangible and direct touch interfaces for manipulating 2d and 3d information on a digital table. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, pp. 77–84. ACM, 2009.

[27] O. Hilliges, D. Baur, and A. Butz. Photohelix: Browsing, sorting and sharing digital photo collections. In *Second Annual IEEE International Workshop on Horizontal Interactive Human-Computer Systems (TABLETOP'07)*, pp. 87–94. IEEE, 2007.

[28] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '94, pp. 452–458. ACM, New York, NY, USA, 1994. doi: 10.1145/191666.191821

[29] K. Hinckley, J. Tullio, R. Pausch, D. Proffitt, and N. Kassell. Usability analysis of 3d rotation techniques. In *Proceedings of the 10th annual ACM symposium on User interface software and technology*, pp. 1–10, 1997.

[30] H. Ishii and B. Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*, pp. 234–241. ACM, 1997.

[31] P. Issartel, L. Besançon, T. Isenberg, and M. Ammi. A tangible volume for portable 3d interaction. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*, pp. 215–220. IEEE, 2016.

[32] B. Jackson, T. Y. Lau, D. Schroeder, K. C. Toussaint, and D. F. Keefe. A lightweight tangible 3d interface for interactive visualization of thin fiber structures. *IEEE transactions on visualization and computer graphics*, 19(12):2802–2809, 2013.

[33] R. J. Jacob, L. E. Sibert, D. C. McFarlane, and M. P. Mullen Jr. Integrality and separability of input devices. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 1(1):3–26, 1994.

[34] S. Jordà, M. Kaltenbrunner, G. Geiger, and R. Bencina. The reactable. In *ICMC*, 2005.

[35] H. Kato, M. Billinghurst, I. Poupyrev, K. Imamoto, and K. Tachibana. Virtual object manipulation on a table-top ar environment. In *Proceedings IEEE and ACM International Symposium on Augmented Reality (ISAR 2000)*, pp. 111–119. Ieee, 2000.

[36] M. Krichenbauer, G. Yamamoto, T. Taketom, C. Sandor, and H. Kato. Augmented reality versus virtual reality for 3d object manipulation. *IEEE transactions on visualization and computer graphics*, 24(2):1038–1048, 2017.

[37] G. A. Lee, G. J. Kim, and M. Billinghurst. Interaction design for tangible augmented reality applications. In *Emerging Technologies of Augmented Reality: Interfaces and Design*, pp. 261–282. IGI Global, 2007.

[38] G. A. Lee, C. Nelles, M. Billinghurst, M. Billinghurst, and G. J. Kim. Immersive authoring of tangible augmented reality applications. In *Proceedings of the 3rd IEEE/ACM international Symposium on Mixed and Augmented Reality*, pp. 172–181. IEEE Computer Society, 2004.

[39] T. Lee and T. Hollerer. Handy ar: Markerless inspection of augmented reality objects using fingertip tracking. In *2007 11th IEEE International Symposium on Wearable Computers*, pp. 83–90. IEEE, 2007.

[40] J. Liang and M. Green. Jdcad: A highly interactive 3d modeling system. *Computers & graphics*, 18(4):499–506, 1994.

[41] T. Louis and F. Berard. Superiority of a handheld perspective-coupled display in isomorphic docking performances. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces*, pp. 72–81. ACM, 2017.

[42] B. MacIntyre. Authoring 3d mixed reality experiences: Managing the relationship between the physical and virtual worlds. *At ACM SIGGRAPH and Eurographics Campfire: Production Process of 3D Computer Graphics Applications-Structures, Roles and Tools, Snowbird, UT*, pp. 1–5, 2002.

[43] D. P. Mapes and J. M. Moshell. A two-handed interface for object manipulation in virtual environments. *Presence: Teleoperators & Virtual Environments*, 4(4):403–416, 1995.

[44] M. R. Mine. Virtual environment interaction techniques. *UNC Chapel Hill CS Dept*, 1995.

[45] S. Miyafuji, T. Sato, Z. Li, and H. Koike. Qoom: An interactive omnidirectional ball display. In *Proceedings of the 30th annual acm symposium on user interface software and technology*, pp. 599–609. ACM, 2017.

[46] S. Miyafuji, M. Sugasaki, and H. Koike. Ballumiere: Real-time tracking and spherical projection for high-speed moving balls. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces*, pp. 33–37. ACM, 2016.

[47] E. W. Pedersen and K. Hornbæk. Tangible bots: interaction with active tangibles in tabletop interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2975–2984. ACM, 2011.

[48] G. Perelman, M. Serrano, M. Raynal, C. Picard, M. Derras, and E. Dubois. The roly-poly mouse: Designing a rolling input device unifying 2d and 3d interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 327–336. ACM, 2015.

[49] J. S. Pierce, B. C. Stearns, and R. Pausch. Voodoo dolls: seamless interaction at multiple scales in virtual environments. In *Proceedings of the 1999 symposium on Interactive 3D graphics*, pp. 141–145. ACM, 1999.

[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*, pp. 79–80. Citeseer, 1996.

[51] I. Poupyrev, T. Ichikawa, S. Weghorst, and M. Billinghurst. Egocentric object manipulation in virtual environments: empirical evaluation of interaction techniques. In *Computer graphics forum*, vol. 17, pp. 41–52. Wiley Online Library, 1998.

[52] I. Poupyrev, D. S. Tan, M. Billinghurst, H. Kato, H. Regenbrecht, and N. Tetsutani. Tiles: A mixed reality authoring interface. In *Interact*, vol. 1, pp. 334–341, 2001.

[53] I. Poupyrev, S. Weghorst, and S. Fels. Non-isomorphic 3d rotational

techniques. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pp. 540–547, 2000.

[54] S. Reed, O. Kreylos, S. Hsi, L. Kellogg, G. Schladow, M. Yikilmaz, H. Segale, J. Silverman, S. Yalowitz, and E. Sato. Shaping watersheds exhibit: An interactive, augmented reality sandbox for advancing earth science education. In *AGU Fall Meeting Abstracts*, 2014.

[55] H. Saidi, M. Serrano, P. Irani, and E. Dubois. Tdome: a touch-enabled 6dof interactive device for multi-display environments. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 5892–5904. ACM, 2017.

[56] H. Saidi, M. Serrano, P. Irani, C. Hurter, and E. Dubois. On-body tangible interaction: using the body to support tangible manipulations for immersive visualization. In *Proceedings of the 30th Conference on l'Interaction Homme-Machine*, pp. 1–11. ACM, 2018.

[57] D. Schmalstieg, L. M. Encarnação, and Z. Szalavári. Using transparent props for interaction with the virtual table. *SI3D*, 99:147–153, 1999.

[58] R. Schmidt, K. Singh, and R. Balakrishnan. Sketching and composing widgets for 3d manipulation. In *Computer Graphics Forum*, vol. 27, pp. 301–310. Wiley Online Library, 2008.

[59] B. Shneiderman. Direct manipulation: A step beyond programming languages. In *ACM SIGSOC Bulletin*, vol. 13, p. 143. ACM, 1981.

[60] P. Song, W. B. Goh, W. Hutama, C.-W. Fu, and X. Liu. A handle bar metaphor for virtual object manipulation with mid-air interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1297–1306. ACM, 2012.

[61] A. Steed. Towards a general model for selection in virtual environments. In *3D User Interfaces (3DUI'06)*, pp. 103–110. IEEE, 2006.

[62] R. Stoakley, M. J. Conway, and R. Pausch. Virtual reality on a wim: interactive worlds in miniature. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 265–272, 1995.

[63] C. Ware and S. Osborne. Exploration and virtual camera control in virtual three dimensional environments. *ACM SIGGRAPH computer graphics*, 24(2):175–183, 1990.

[64] C. Ware and J. Rose. Rotating virtual objects with real handles. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 6(2):162–180, 1999.

[65] F. Zhou, H. B.-L. Duh, and M. Billinghurst. Trends in augmented reality tracking, interaction and display: A review of ten years of ismar. In *Proceedings of the 7th IEEE/ACM international symposium on mixed and augmented reality*, pp. 193–202. IEEE Computer Society, 2008.