TabletInVR: Exploring the Design Space for Using a Multi-Touch Tablet in Virtual Reality

Hemant Bhaskar Surale
Cheriton School of Computer Science
University of Waterloo
Waterloo, Canada
hsurale@uwaterloo.ca

Aakar Gupta
Cheriton School of Computer Science
University of Waterloo
Waterloo, Canada
aakar.gupta@uwaterloo.ca

Mark Hancock
University of Waterloo
Waterloo, Canada
mark.hancock@uwaterloo.ca

Daniel Vogel
Cheriton School of Computer Science
University of Waterloo
Waterloo, Canada
dvogel@uwaterloo.ca

ABSTRACT

Complex virtual reality (VR) tasks, like 3D solid modelling, are challenging with standard input controllers. We propose exploiting the affordances and input capabilities when using a 3D-tracked multi-touch tablet in an immersive VR environment. Observations gained during semi-structured interviews with general users, and those experienced with 3D software, are used to define a set of design dimensions and guidelines. These are used to develop a vocabulary of interaction techniques to demonstrate how a tablet's precise touch input capability, physical shape, metaphorical associations, and natural compatibility with barehand mid-air input can be used in VR. For example, transforming objects with touch input, "cutting" objects by using the tablet as a physical "knife", navigating in 3D by using the tablet as a viewport, and triggering commands by interleaving bare-hand input around the tablet. Key aspects of the vocabulary are evaluated with users, with results validating the approach.

KEYWORDS

interaction techniques, virtual reality, touch interaction

ACM Reference Format:
Hemant Bhaskar Surale, Aakar Gupta, Mark Hancock, and Daniel Vogel. 2019. TabletInVR: Exploring the Design Space for Using a Multi-Touch Tablet in Virtual Reality. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3290605.3300243

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-1-4503-5970-2/19/05.
https://doi.org/10.1145/3290605.3300243

1 INTRODUCTION

While virtual reality (VR) has been around in various forms since at least the 1960s (e.g., [54]), advances in display technology have sparked a new interest from both researchers and the public. There are clear advantages to virtual reality, like the ability to look and move around in an immersive 3D environment. Yet, VR interaction is challenging due to limited tactile feedback, poor input precision when drawing [2], and lack of a consistent interaction vocabulary. Past research has introduced methods for haptic feedback [5, 16, 37, 52, 60], techniques to increase precision [43], and more standardized control schemes [49]. In our work, we leverage the familiarity and ubiquity of multi-touch tablets as a means of interacting with 3D content in a VR world.

We introduce a "TabletInVR" design space combining a 3D-tracked tablet with mid-air barehand gestures, which we demonstrate in an example interaction vocabulary for 3D modelling.

Exploring VR interaction in the context of 3D modelling is particularly compelling because the task *should* be a good fit for VR, but in practice, supporting the many required operations is challenging (e.g. object creation, selection, transformation; world navigation; copy, paste, undo, etc.). Although past research has considered the use of 2D surfaces in VR, this has focused on 3D-tracked props without real multi-touch input [34, 35, 44], or using multi-touch tablets for transforming 3D objects without exploiting 3D tablet tracking [15, 46].

Our work combines the affordances of a 3D-tracked tablet with the input capabilities of its multi-touch surface. We advocate that the tablet's precise touch input capability, physical shape, metaphorical associations, and natural compatibility with barehand, mid-air input can be effectively used in VR. Interactions involving precise mutli-touch input could begin on the tablet followed by coarse hand gestures in VR, or tablet input could be used to transform objects or navigate the world in a familiar mutli-touch way. This suggests interesting aspects when combining these two modalities. Interactions can leverage physical qualities like the 2D tablet input providing a continuous tactile sensation and a mid-air gesture enabling free movement in space. Interactions spanning the tablet and mid-air gestures could more effectively exploit bimanual input, since bimanual multi-touch input is known to work well [21, 58], but pure bimanual mid-air gestures are less reliable in VR [28, 50].

Our primary contribution is the definition and exploration of a design space of using a multi-touch tablet in VR.

In a formative study, we asked participants familiar with 3D software to envision how they would perform standard 3D modelling tasks using a tablet in VR.

Based on observations of behavioural patterns and proposed features, we mapped out a design space with eleven dimensions (e.g., 'physical vs. non-physical', 'direct vs. indirect', and 'discrete vs. continuous') and developed a vocabulary of interactions (e.g., 'two-finger drag to translate an object', 'five-finger drag to navigate', and 'swipe-in to delete the object'). Lastly, we validated our system similar to Arora et al.[1], where participants created 3D models to test the design space.

2 BACKGROUND AND RELATED WORK

In this section, we summarize prior investigations in the domain of mixed reality environments, using props as tablet, touch input, tablet and pen, and hand gesture input, while focusing on 3D solid modeling.

Tablet and Pen in Mixed Reality. One common approach found in past work is to use a tablet with a pen as an input device in VR. Aspin et al. [4] explored the use of a 3D tracked tablet and stylus in a CAVE-like system for navigation and exploration of small, complex 3D structures. Bowman et al. [11] explored using a tablet and stylus in VR for the assessment of building structures. Billinghurst et al. [8] used a pen operated pressure sensitive pad to support content creation in a virtual environment. Similarly, Bornik et al. [10] used a 6-DOF tracked pen with a tablet to view and manipulate medical data, Reitmayr and Schmalstieg [45] explored the use of a pen and a tablet-like pad (both are props tracked using markers) for a collaboration task, and Sareika et al. [48] investigated bimanual interactions for urban planning using a pen and tablet. Keefe et al. [32] explored precise mid-air strokes using a haptic-aided input technique for 3D sketching, and Arora et al. [1, 2] investigated the impact of the lack of a physical surface on drawing inaccuracies. Their work explored both 3D sketching in augmented reality (AR) using a mid-air pen-based drawing and 2D surface sketching . More recently, Aslan et al. [3] conducted a series of studies to gauge the potential of pen and mid-air input and noted that mid-air input should complement pen and touch-enabled tablets. However, compared to modern high-fidelity multi-touch tablets, pen input is essentially limited. It does not take the full advantages of direct multi-touch finger input.

Albeit sketching using a pen is not our focus, we look at these results through the lens of interaction design. A few lessons to learn before we address 3D modeling. For instance, tablet's continuous input on tablet surface; drawing in mid-air with the help of the tablet, orientating tablet surface in arbitrary plane, and so forth.

Tabletop and Hand Gestures in Mixed Reality. Benko et al. [6, 7] explored the interaction space combining a tabletop and hand gestures in a partially immersive environment. Marquardt et al.'s *continuous interaction space* [36] and Mockup-Builder [17] both demonstrated these sorts of 2.5D interactions. They start on a planar surface and continue in mid-air, something highly promising for 3D modeling applications. However, a tabletop surface does not have the same flexibility as a hand-held tablet, such as orientation tracking, mobility, a mid-sized display, and so forth. Being able to carry the tablet around facilitates interactions without being physically constrained to a certain position.

Tablet as a Prop in Mixed Reality. Another way to provide 2D input in VR is to use a prop like a tablet, without a touch sensor. For instance, Linderman et al. [35] demonstrated the use of a passive-haptic paddle as a 2D input device for widget selection in VR, Poupyrev et al. [44] used it for text-based applications (note-taking, text input, and annotation using physical pen as a prop), and Szalavari et al. [55] used it for 3D modeling application. However, none of the past efforts have explored the simultaneous use of multi-touch tablets with hand gestures like in our work.

Tablet touch in Mixed Reality. Wang and Lindeman [59] presented an AR environment consisting of a semi-transparent HMD (Head Mounted Display), a wand in the right hand, and a multi-touch tablet mounted on the left forearm. The interface enabled looking at the virtual environment, as well as seeing the tablet mounted on the non-dominant hand of the user. However, tablet touch interaction was cumbersome since the wand had to be held somewhere other than the right hand temporarily (in the left hand, or between the legs). Kim et al. [33] explored a scaled-down locomotion that allows a user to travel in a virtual world as their fingers slide on a multi-touch surface. However, finger motions were not precisely detected and only two finger touch was investigated . In contrast, our system adds more expressivity by tracking two or more touch points [61], utilizing device orientation [47] to navigate in arbitrary plane and using mid-air hand gestures [14, 42].

3 FORMATIVE STUDY

The goal of this formative study is to gain insights into how people envision using a physical tablet in a VR environment, using the context of a 3D modelling application. Our approach is similar to work by Hinckley et al., who observed how people used physical paper and notebooks to inform new design spaces for combining touch and stylus input [29] and stylus grip sensing [27]. Like those works, the observations from this formative study are later used to build a design space and an example interaction vocabulary, and we also conduct a preliminary user study to validate our design space through the example interaction vocabulary.

Participants. Ten people (7 male, 3 female, ages 22–26) participated. Three were architecture students experienced with Revit, Fusion360, Sketchup, and SolidWorks software; two were mechanical engineering students experienced with SolidWorks and Fusion360; and two were amateur users with some experience using software like 123D, Blender, and Sketchup. Five participants had experience with VR.

Table 1: List of the tasks and corresponding actions suggested by formitive study participants.

Task	Subtask	Suggested Actions
Create	Primitive	Draw on tablet and extrude or push away, menu buttons
	Clone	Grab longer and drag
Select	Object	Grab, tap on tablet
	Face	Tap on object
	Group	Non dominant grab, two finger, lasso on tablet
Transform	Rotate	Rotate hand
	Translate	Move hand
	Scale	Pinch to zoom, distance between two hands
Modify shape	Slice	Menu buttons, slice with a hand, tablet to slice
	Extrude	Draw 3D path, pinch face
Modify texture	Colouring	Menu buttons

Procedure. Participants were told to imagine how they would use a multi-touch tablet together with mid-air barehand gestures in a fully immersive 3D environment containing 3D objects. Like past work [27, 29], participants were asked to act out specific tasks. In our case, these were basic 3D modelling operations (see Table 1). Each participant used a multi-touch tablet (which was not turned on). A chair was provided, but the decision to sit or stand was left to the participant. Two small and one medium sized cardboard cubes were placed around the user, some within arm's reach, and some beyond. These cubes helped them visualize an object to create or manipulate without wearing an HMD.

We asked each participant to imagine and act out 3D object creation, selection, manipulation, and annotation. While performing the tasks, they simultaneously explained their envisioned system using a think-aloud protocol. This included the steps they took, and their opinions about important considerations and choices they made. Observations were recorded by the experimenter as written notes.

Observations

We observed participants' behaviour and analyzed notes using affinity diagramming to reason about the role of a tablet in VR. Our design space is a manifestation of the following seven core observations.

Delegation of Tasks:

O1. *Granular and coarse actions*: participants preferred using mid-air hand gestures for coarse actions, followed by input on the tablet for finer control, "I'd grab an object and then use the tablet to rotate it." [P3]

O2. *Near and far actions*: instead of navigating to a distant object, beyond arm's reach, participants preferred indirect object selection using the tablet. For instance, the tablet's screen could depict a birds-eye view [51, 59], where a tap on the tablet selects an object. Participants also suggested treating the tablet as a remote control, so they could raycast to select. However, to select objects within arm's reach, they preferred to reach out and grab with their hands.

Tablet Properties:

O3. *Tablet as interface*: participants suggested using menu buttons (2D) on the tablet to create objects, invoke commands, and select modes. Although they utilized a mixture of mid-air hand and touch-based gestures, most tasks were initiated on the tablet with a tap of a button. Tracking the tablet orientation and position creates novel precision-focused interactions. For example, to translate an object, a user can tap on the tablet to select it, and then drag with their fingers to translate while adjusting the translation axis through the tablet's orientation.

O4. *Tablet as a tool*: Despite not being common, a few participants used the tablet to define a slicing plane, and some used their dominant hand for slicing an object (like a knife). this behaviour was from a fruit ninja game, where players use a sword to cut through fruits. Other participants used bare hands to slice an object, but were skeptical about accuracy and unsure it was a suitable operation.

The physical form of the tablet affords a variety of operations when tracked and rendered virtually in VR. It can be made to resemble a knife, a tray, a rectangular block, a ruler, a storage unit, among other physical forms. The plane of a tablet can be aligned with the face of an arbitrary 3D object to extrude, color, or even delete it. A corner of a tablet can be used as a pointer, which can be used to select objects.

O5. *Haptic feedback*: mid-air hand gestures seemed suitable for discrete interaction and touch-based interactions were favoured for continuous manipulation. This behaviour appeared to be linked with the demand for haptic feedback and the perceived precision requirements of the task. For instance, participants suggested hand gestures to grab an object and a pinch-to-zoom gesture on the tablet to scale it. Moreover, prior research has shown that haptic sensations in VR can greatly improve the user experience [16, 30, 60].

Symbolic input and UI interactions (e.g., buttons, menus, etc.) can also benefit from having a physical, tactile surface. The tablet can act as an arbitrary UI (e.g., to annotate objects or select modelling operations), and the tactile feedback can improve typing speed when compared to mid-air typing without haptic feedback [22].

General Observations:

O6. *Occlusion avoidance*: participants felt that using a tablet for continuous manipulation tasks made more sense than using hand gestures, as it avoids occluding the object of interest and requires minimal efforts, "[...] and my hands will not even occlude the object." [P1].

O7. *Sit vs. Stand*: all participants preferred sitting, except one who demonstrated a willingness to stand, "I can stand if I have to look at the cube from the top side." [P2], but still opted to sit throughout the exercise.

4 DESIGN SPACE

Following a systematic approach, we consider these observations (*O*1 - *O*7), depicting user behaviour, to build the design space, followed by designing the interaction vocabulary. For each candidate dimension, in design space, we pose intriguing questions that instigate design considerations. Such considerations would help interaction designers assign different roles to the tablet. Note that our focus is on building a design space and interactions for using a tablet in VR. We do not contest to investigate 2D input or hand gesture input in isolation [9, 31, 61] or passive 2D input in VR, as it has been studied elsewhere [18, 35].

Informed by our set of high-level observations from the formative study along with past research, we shape a design space. Recall that we asked participants to envision interactions for three settings, so, we assemble these interactions together with corresponding dimensions in each setting to bring out novel and rational interactions. These dimensions are essentially the lenses thorough which we can envision the possibilities of the TabletInVR concept.

Design Dimensions

Design dimensions are the core components of interaction space, where each interaction we envision is composed of one or more of the following design dimensions.

Table 2: Tablet vs. Mid-air properties.

In VR / **Properties**	Tablet touchscreen	Mid-air tablet	Mid-air hand
Precision	High	Low	Low
Input space	2D (3DOF)	3D (6DOF)	3D (high DOF)
Tactile feedback	Yes	No	No
UI	Familiar (WIMP)	No (tilt)	No (gesture)
Midas touch	No	Yes	Yes

D1. *Tablet vs. mid-air properties*: Table 2 describes different properties of tablet and mid-air interactions in VR. Tablet in VR could be mutually beneficial given high precision input space on 2D surface. However, when does high precision input is essential? Is 3DOF input adequate? In VR, which interactions need tactile feedback? or UI?

D2. *Non-dominant (ND) vs. dominant (D) hand assignment*: Participants used their dominant hand while using the tablet as a tool (O4). How the ND and D hand roles are defined based on the use of the tablet in VR?

D3. *Sit vs. stand*: Body posture can have an impact on fatigue and on the interaction experience in general. As pointed out in O7, only one participant was willing to stand, and this depended on the task. What tasks are suited to sitting vs. standing? Can VR provide the flexibility to either stand or sit irrespective of the task?

D4. *Attention to device vs. scene*: Recently, Yan et al. [63] found that, compared to eyes-engaged, the eyes-free approach is significantly faster, provides satisfying accuracy, and introduces less fatigue and sickness. Can interaction spaces be divided into different regions, either on the tablet or around the user, to guide attention and leverage the benefits of eyes-free interaction?

D5. *Unimodal vs. multimodal*: There are many modes of interaction for a tablet in VR. For instance, combined mid-air hand and touch gestures, touch-only, unimanual or bimanual hand gestures, tablet orientation and touch, and so forth. How and when can these modes be applied to reduce fatigue or to improve accuracy, and in general, reduce user frustration?

D6. *Unimanual vs. bimanual*: Past research has explored the benefits of bimanual interaction [13, 23, 56]; however, bimanual interaction may not be suitable for every task, for example, grabbing a virtual object using only the dominant hand. Should a task be performed using either one or both hands? Does it improve the task completion time? Modern VR devices are equipped with reasonably accurate hand-gesture recognizers. Mid-air barehand gestures along with touch gestures enable unique workflows. For instance, a pinch gesture could be used to select an object, followed by a pinch-to-zoom on the tablet. A long pinch could be used to create a ghost copy of an object, followed by a two-finger rotation gesture on the tablet.

D7. *Environment reality vs. virtuality*: There exist multiple ways to provide different levels of visual feedback. For instance, with a standard tablet, the 3D world can exist only in the confined window of a tablet screen. Similarly, when the tablet is tracked to create a viewport, the world can be seen through a tablet screen; however, the virtual objects are physically stuck to the real environment, like in augmented reality. The tablet acts as a portal to the virtual world around the user [24]. Furthermore, when a tablet is used in VR, a portal could let a user view the real world while being in VR [38], creating a 'portal to reality'. While wearing an HMD, the ability to be aware of one's surroundings is essential. With the tablet in hand, a user can peek into reality whenever desired. Prior research has explored using a flat surface to create a viewport [55], however, interacting through a viewport is an unexplored area of research. Further, When transitioning between modes of operation, how and when can an awareness of the real world be provided in the virtual world and vice versa? Can the tablet' s screen used as a portal to and from reality? Does it break the immersion?

D8. *Interleaved vs. simultaneous*: Does an interaction require simultaneous use of both input techniques, such as touch and mid-air hand gestures, or given a task would only one of them would suffice? Is it preferable to use both input techniques in an interleaved fashion, or to solely rely on one of them?

D9. *Discrete vs. continuous input*: While tapping on a virtual object to select it is an example of discrete input, changing the scale of an object is an example of continuous input. However, the mapping from task to input type is not always clear. For instance, consider a relatively complex task of selecting an object from a stack of objects, which is placed beyond arms reach, would we still resort to discrete input or would mixed input be more efficient? What scenarios drive such a mapping? When should a designer opt for discrete input and when should they opt for continuous input?

Furthermore, tablets provide a high-fidelity input space in a low-fidelity virtual environment. We can go beyond taps and clicks to recognize hand-drawn gestures in VR. This would allow users to provide a more advanced form of input. For instance, gesture-based menu invocation [64] and hand-drawn shape recognition could be used to invoke commands, or more traditional forms of input such as pinch-to-zoom and two-finger swipe could be used.

D10. *Direct vs. indirect*: As pointed out in O6, to avoid occlusion, participants used the tablet screen instead of mid-air hand gestures. Similarly, in O2, participants used the tablet screen to select distant objects. These observations hint toward the need for an indirect manipulation technique. Does the interface leverage the full potential of available input methods for direct and indirect tasks?

D11. *Physical vs. non-physical*: O4 highlighted the use of the tablet as an entity which does not necessarily follow physical laws from the real world. We identify this being a crucial factor while assigning roles to the tablet and the user's hands in VR. We try to reason about the possibility of assigning direct interaction with the hands to abide by physical laws of the real world and non-physical interactions using the tablet. For instance, direct tap using a finger might displace a virtual object, while the tablet could pierce through a virtual object to select it. Moreover, could such physical and non-physical interactions lead to a better experience of using tools in VR? Would switching such roles make interactions difficult and unusable? In what cases should direct interactions using hands not follow physical laws?

5 EXAMPLE INTERACTION VOCABULARY

We describe an example interaction vocabulary for using a tablet in VR for the purpose of 3D solid modelling. The interaction techniques are informed by the formative study observations and are constructed to span and illustrate the design space dimensions. Table 3 shows how the design dimensions informed the interactions. For instance, the *Create* interaction is a result of flipping the tablet using the non-dominant hand (D2) and discrete taps (D1, D6, D9) on the back of the tablet using the dominant hand. Each family of interaction techniques are described generally, with specific implementation details from our application provided to make the ideas more concrete.

TabletInVR Prototyping System. How the interaction techniques were implemented was partially influenced by the capabilities of our prototyping system. The application runs in Unity 5.6.1, on a high-end Windows 10 machine (3.6GHz Intel i7 CPU, 1.6GHz GeForce GTX 1080 GPU). The VR HMD is an HTC Vive (1080 × 1200 px per eye, 90Hz refresh, 110° fov). Hand tracking uses a LEAP motion device mounted on the front of the HMD with the interaction engine v1.1.1. The mounting angle and 135° field-of-view of the LEAP camera enables the hand using the tablet to be tracked when the user looks at the tablet. The interaction engine Unity LEAP plug-in displays 3D models of the users hands in VR. It should be noted that LEAP hand tracking is not robust to IR reflection, especially when the tablet is near and when the finger positions are pointing away from the LEAP. As a result, our implementation avoids these kinds of in-air gestures with touch interaction available on the tablet.

The tablet is a 9.7" Samsung Galaxy Tablet S3 (1536 × 2048 px display, 264 ppi), weighing 429g. The 3D position and

Table 3: Interactions used in vocabulary with mappings from TabletInVR design dimensions.

Dimensions \ Interactions	Create	Select, Deselect	Delete	Transform	Modify	Annotate	Navigation	Help	System Menu
D1. Tablet vs. mid-air properties	Tap on tablet	-	Mid-air gesture	Drag on tablet	-	-	-	-	-
D2. Hand assignment	Dominant	-	-	-	Dominant	-	-	-	Non-dominant
D3. Sit vs. stand	-	-	-	-	-	-	Both	-	-
D4. Attention to device vs. scene	-	-	-	-	-	On the device	On the device	On the scene	-
D5. Unimodal vs. multimodal	-	-	-	-	Tablet orientation and touch	Touch	Tablet orientation and touch	-	-
D6. Unimanual vs. bimanual	Bimanual	Both	Unimanual	Bimanual	-	-	Bimanual	Bimanual	-
D7. Environmental reality vs. virtuality	-	-	-	-	-	-	Tablet viewport	-	-
D8. Interleaved vs. simultaneous	-	Interleaved	Interleaved	-	-	-	-	-	-
D9. Discrete vs. continuous	Discrete	-	-	Continuous	-	Discrete	Continuous	-	-
D10. Direct vs. indirect	-	Both	-	-	-	-	-	-	-
D11. Physical vs. non-physical	-	Both	-	-	Non-physical	-	-	-	-

orientation of the tablet is tracked using an HTC Lighthouse tracker. Tracking of this tracker is glitchy when the docking area is facing the Lighthouse [57], however, it did not hinder the usability. The 9.9cm × 4.2cm tracker is screwed into a lightweight aluminum bar, which is attached to the back of the tablet using high-strength hook-and-loop fasteners. The bar is attached such that the tracker extends approximately 4.5cm out beyond one corner. This mounting position enables portrait and landscape orientation when held in one or two hands, and the tablet can be flipped to use the back as a haptic surface. We perform calibration of the virtual tablet model and the physical tablet manually, by adjusting the rotation and translation offsets until they align. Since the tracker is securely fixed to the tablet, this one-time manual calibration is acceptable. Multitouch events registered by the tablet (x and y coordinates of each touch point) are sent in real-time to the server over a high speed WiFi network.

3D Modelling Tasks. Foley et al. [20] provide a fundamental interaction task set, independent of application and hardware, for 3D environments—select, position, orient, path, quantify, and text. Our interaction vocabulary includes these tasks and builds upon them with more advanced interactions: selection, deselection, manipulation (rotate, scale, translate), modify (slice and extrude), creation, deletion, annotation, and so on. In our interaction vocabulary that follows, we demonstrate how a tablet's precise touch input capability, physical shape, metaphorical associations, and compatibility with barehand mid-air input can be used in VR to perform these 3D modelling tasks.

Tablet Viewport

We design several interactions to use the affordance and physical properties enabled by a view of the 3D scene rendered in the HMD's view of the tablet display. In addition, we add a 3D ray emanating out from the centre of the back of the tablet. The combination of the viewport rendering, the ray, and available multi-touch input of the tablet creates a useful direct and indirect interface.

Note that the viewport rendering is only on the front display side of the tablet, and it is hidden when the tablet is being used for other purposes, like rotating, translating or scaling an object, or when using the tablet to slice objects. This helps the user maintain focus on the object(s) being transformed (D4). During navigation, only this viewport rendering is visible with the virtual world made uniformly black. This enables the viewport rendering to function as a small view of the 3D scene during navigation, which is less likely to induce motion sickness [19, 41].

Creation

To select an object for creation, a user flips the tablet over with their non-dominant hand, browses a list of objects using

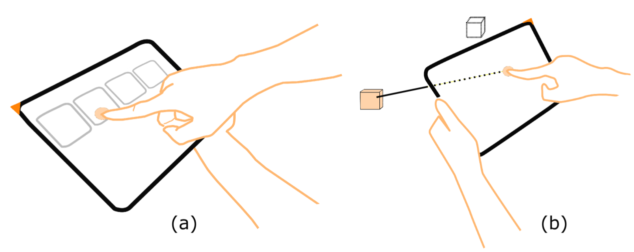

Figure 1: Creation. (a) Flip the tablet, select the object for creation, (b) Tap on tablet viewport to create.

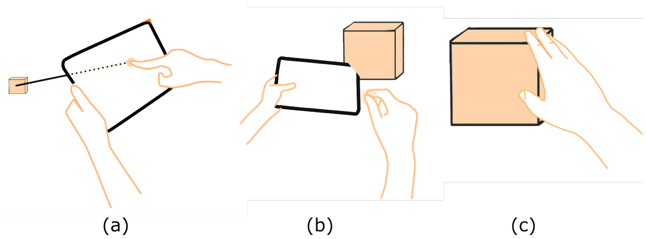

Figure 2: Object selection. (a) Tap on tablet viewport, (b) Pierce tablet corner in the object and pinch, (c) Tap on the face of the object.

a scrolling list, and taps on one to select it (O3, D6, D9). The selected object appears as a 3D icon near the top of the tablet, indicating creation mode is active and what object has been selected for creation (see Figure 1).

When an object is selected, there are two ways to create it in the scene. First, the user can *remotely* create objects by pointing the ray from the back of the tablet at the grid on the ground plane. A tap on the tablet screen creates an object at that point on the grid (O5). Second, the user can create objects in *mid-air* (50cm in front of them) by pinching the thumb and index finger of their dominant hand while holding the tablet in their non-dominant hand (D1, D2).

Multiple objects can be created by repeating either a remote tap or mid-air pinch, and a different object can be created by flipping the menu and selecting a different object. To exit creation mode, a "swipe-in" movement is performed using the dominant hand just over the surface of the tablet. This follows the affordance of brushing off the icon of the creation object (O1, D9). We use the same gesture for deleting a selected object, explained later.

Our application supports primitive-shape creation (cube, cylinder, sphere, capsule) and Minecraft-style [40] blocks.

Selection (and Deselection)

In VR, selection methods differ based on how far away the object of interest is (O2, O4). So, we employ three different selection methods that take advantage of the tablet's form in conjunction with hand-tracking (Figure 2).

First, we use the tablet viewport for selecting a distant object, usually beyond arm's reach and within sight (O2). The

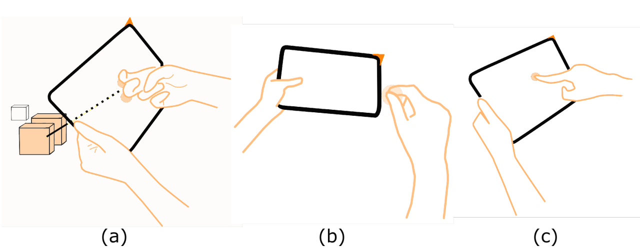

Figure 3: (a) Knuckle for multiple object selection, (b) Pinch to deselect, (c) Tap on viewport to deselect.

user points the tablet's ray at a distant object and taps on the tablet screen to select it (O3, D10). Second, a corner pointer with a bright yellow highlight at the top right corner of the tablet can be used to select an object by first piercing through the object with the corner, and then using a dominant-hand 'pinch' gesture (O2, O4, D6, D8, D11). Third, a user can also select an object by tapping with their dominant hand on the face of the object (O2, D6, D10).

In order to select multiple objects, a knuckle hand posture (see Figure 3, described in TapSense [25] and in the mode-switching study by Surale et al.[53]) is used along with one of the selection methods. Selected objects are highlighted using custom shaders (orange color). When the tablet is piercing two adjacently placed objects, only the object enclosing the corner pointer will be highlighted yellow (O6). Highlighting is used to indicate the hover state (O2, O4, D10, D11). This makes the corner pointer a precise object selection method, especially in case of a cluttered scene. Objects can be individually deselected by selecting them again using any of the techniques, and all objects are deselected when selecting "nowhere" with the corner pointer or tablet viewport, or selecting another object (Figure 3 (b-c)).

Deletion

Deletion follows selection. To delete an existing object, select it and perform a "swipe-in" movement using the dominant hand just over the surface of the tablet (see Figure 4 (a)).

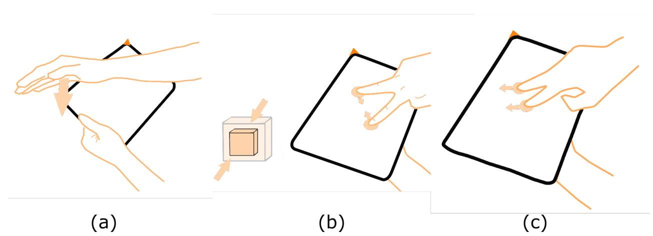

Figure 4: (a) Swipe-in to delete the object, (b) Two finger scale, (c) Two finger drag to translate.

Rotate, Scale, and Translate

Rotate, scale, and translate transformations (Figure 4(b-c)) follow selection, and can be performed simultaneously. Once an object is selected, orienting the tablet fixes an axis and plane of transformation (Figure 5) (O1, O3, O5, O6, D9). The tablet orientation has to be maintained until the end of the transformation. Two-finger touch on the tablet starts transformation, releasing the contact disengages. Users can perform transformations with 9DOF.

Modify

The shape of an object can be modified by either slicing the object in an arbitrary plane or selecting a face for extrusion. Both operations follow selection. A user can slice an object by placing the tablet through the object with their non-dominant hand to determine a slicing plane and then using a dominant-hand pinch gesture to trigger the slice (D2, D5, D11). The half of the object above the tablet's screen will be removed, and the remaining portion of the object will be kept (Figure 6 (a)). To extrude, choose a face of an object by orienting the tablet in one of the three orientations (similar to Figure 5). Once the desired face is selected, two finger horizontal drag on the tablet extrudes the face (inward or outward). Our application supports extrusion with cube(s).

Text Annotation

Having a physical tablet has the major benefit of providing a means for text input (O5). To create an annotation, we leverage the creation techniques and add 'text' to the scrolling list of available objects. Thus a user can place it in the scene remotely or in mid-air, similar to the Virtual Notepad system [44]. Annotations are interactive objects and can be selected or repositioned. When selected, a keyboard will appear on the tablet for typing and 'Enter' is used to commit the changes (O3, O5). While editing, a textbox will appear just above the keyboard showing the current annotation text. This helps maintain the focus on typing without needing to look at the annotation object directly.

Navigation

Simulator sickness or motion sickness is a well-known issue in virtual reality, and navigation without physical movement can exacerbate the problem. One effective way to mitigate this issue is to limit the user view. For instance, Fernandes et al.[19] used varying sized vignettes to limit the visual input to the user resulting in reduced motion sickness. An extreme version of this is recommended by Oculus [41], by fading the scene entirely to black.

We employ a similar approach. Five-finger touch on the tablet initiates navigation (see Figure 6 (b)), and while navigating, the scene quickly fades to black, except for the tablet and viewport. As a result, it is possible to see through the tablet screen to view the scene. The moment navigation stops, by lifting one's fingers off the tablet, the scene is brought back to full visibility with a 3-second fade (O3, D4).

Zoom-in/out, rotate, and drag gestures are used for navigation. To rotate the view, touch the tablet with five fingers and rotate the wrist either to the left or to the right. Fiver-finger drag will intiate a move along the tablet plane (orientation), which can be adjusted with the non-dominant hand (O5, O6, D6, D7). Like transformation, navigation uses two points, the mid-point of the five-finger touch and the first point of contact to enable navigation. Note that five-finger drag moves the person in the opposite direction of the drag, which has the effect of the view rendered on the tablet moving in the same direction as the fingers. The five-finger zoom in/out gesture navigates from the initial position to the forward/backward direction pointed by the tablet, respectively (D5). In our application, rotation rotates the user view around the up axis pointing toward the sky. Note tablet orientation makes no difference for scene rotation.

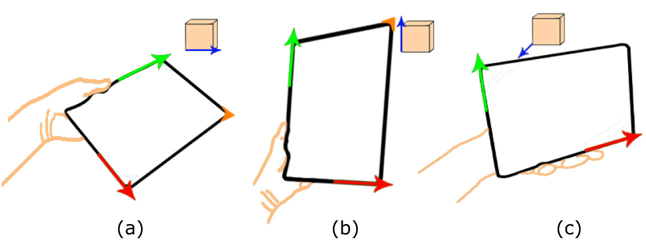

Figure 5: Select axis of transformation using the orientation of the tablet. (a) Facing up to select the x-axis, (c) Portrait vertical to select the y-axis, and (c) Landscape left to select the z-axis.

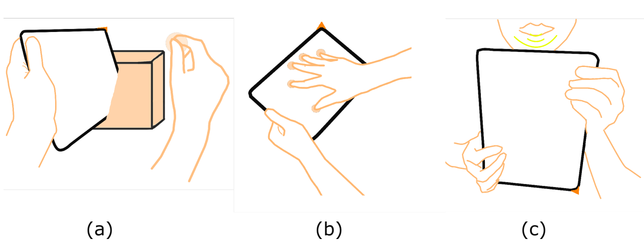

Figure 6: (a) Slicing an object, (b) Five-finger touch to navigate, (c) Speak to tablet and ask for help.

Seeking Help

To request help, the user can old the tablet with two hands (D6) up to their chin (see Figure 6 (c)) and query into the mic (O4). Voice recognition on the tablet responds to the query. Here, a metaphor of a person thinking while holding a writing pad against the chin is used. A quick help video is played a meter in front when the distance between the HMD (D4) and the tablet is within range (≈ 10cm). The video stays in the view as long as the gesture is maintained.

System Menu

Butterworth et al.[12] demonstrated the use of system menu in early work on VR 3D modelling for operations such as undo, redo, cut, copy, and paste. In our system, hold the tablet in the dominant hand to access the system menu (D2). Multiple options are available; For instance, share, clear, exit.

6 USER EVALUATION

Our evaluation protocol and the goals of the user evaluation are similar to Arora et al.'s investigation [1]. We focus on overall usability of the system to replicate a predefined target model, understand user workflow, and analyze user feedback to spot shortcomings. We also ask participants to use our

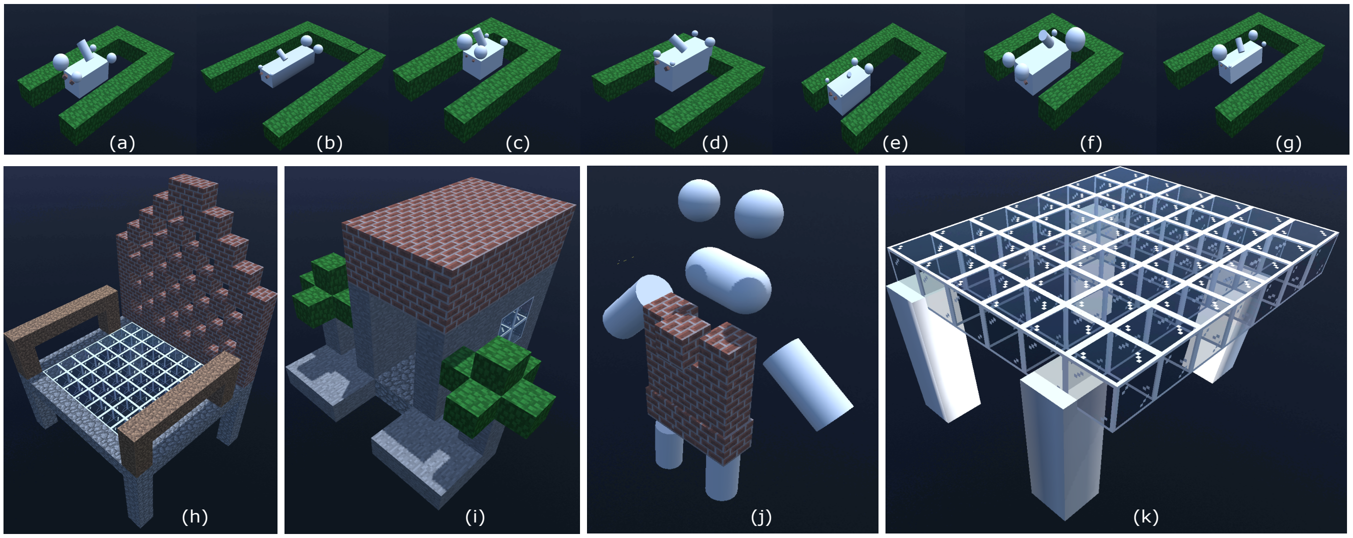

Figure 7: Sample results from 'replication' and 'freeform exploration' task. (a) Target Model, (b-g) Participant's replication (*P*1-*P*6), (h-k) Participant's creations in 'freeform exploration' (*P*1-*P*4)

system to create a model purely out of their imagination. At the end of the study, participants filled out a post-experiment questionnaire indicating their experience with individual features of the system.

Participants. Six people (all male, ages 19-34) participated in our study. They were experienced with Fusion360 (*P*1, *P*2), SolidWorks (*P*2), and other 3D modeling (*P*3-6) tools. All received $15 for successful completion of the study.

Procedure. The study had three parts which took approximately 90 minutes in total to complete:

Part 1: Training (20-30 minutes). Participants were introduced to the system and how to use it.

Then, the experimenter demonstrated and simultaneously explained five main features of the system: create and delete, select/deselect, transform, navigate, and modify. Participants practiced using the main features until they felt confident.

Part 2: Replication Task (20-30 minutes). The purpose was to exercise all the primary features. Essentially, testing the overall usability of the system by making the participants replicate the target model in half an hour. The target model is a predefined spatial arrangement of specific 3D objects placed on the floor as shown in Figure 7 (a). Replication includes completing a set of tasks in any order: 1) Create a cube on the floor using the grid; 2) Extrude it; 3) Create four spheres near the top corners of the cube in decreasing order of scale; 4) Create and place a cylinder on the centre of the top face of the cube; 5) Rotate the cylinder by -45°; 6) Create two 'brick' blocks, scale them down by more than 50% and place it on the front face of the cube; 7) Rotate these blocks by 45°; 8) Create a fence around the cube using 'grass' blocks. Before starting to replicate, participants were asked to familiarize themselves with the target model. The target model was always visible at their front-left side (tilted 45° facing the participant) as a reference. Note they were not required to match exact dimensions of the target object.

Part 3: Freeform Exploration (20-30 minutes). Participants were told to explore the system on their own to make their own creation, and were allowed to search the internet for inspiration. Participants were told to use their preferred features to create the 3D model they imagined. After completing the session, they filled out a post-experiment questionnaire.

Participants were encouraged to take breaks between each part of the study and to notify the experimenter if they were feeling nauseated. However, none of the participants reported feeling discomfort. Except *P*2, all the participants preferred sitting throughout the study.

Results and Qualitative Feedback

All participants successfully replicated the target model within the specified time limit (Figure 7 (b-g)). During the freeform exploration, participants created a chair (*P*6), an android (*P*4), a tree (*P*1), houses (*P*1, *P*5), and a glass table (*P*3) (Figure 7 (h-k)).

Overall, the system was perceived to be useful and interesting: *P*5 noted, "Most of the features, gestures were very intuitive and easy to follow", *P*6 noted, "[...] was an amazing experience, it really felt like we are interacting with the real world objects." Additionally, we analyze user feedback to understand the strengths and the limits of our system.

Create, delete, and modify. Observations and comments from participants indicate that creation, deletion, and modify were intuitive interactions. Except *P*2 (“Sometimes it created objects despite not being intended”), most of the participants could hold the tablet without accidentally touching the tablet screen; however, *P*2 had difficulty holding the tablet in a way that it would not cause unintended taps. To mitigate such problems, we could ignore the touch points near the grip [27] or a provide a longer handle on the left side of the tablet. Also, using design dimension D1, we can assign the trigger to a mid-air pinch [26] or fist gesture [49], and using dimension D2, non-dominant hand touch events can directly be discarded when near the tablet.

Select and deselect. The majority of participants found selection and deselection easy to understand and did master it quickly. Also, a few participants preferred raycast over corner selection. *P*1 felt corner selection was “weird”. All participants found orienting the tablet to select the face of the cube for extrusion to be useful and reported positively. However, multi-object selection using the knuckle hand posture received mixed reviews. Except *P*2 and *P*5, participants reported it to be hard to use. They felt rotating the dominant-hand wrist to be tiresome. Using dimension D6, a dominant hand touch can trigger selection, while a non-dominant hand touch on the tablet can be used to switch between single-object or multi-object modes. Using the non-dominant hand for mode selection while holding the tablet has been effective in prior work [21, 58].

Transform (Scale, Rotate, and Translate). Recall that all three transformations can be performed simultaneously. This approach is similar to many pre-existing tablet applications like maps and image viewers, but for controlling object transformations, participants expressed mixed reviews. Participants found it easy to transform (rotate, scale, and translate) an object, but maintaining the distance between the fingers while rotating proved challenging, as finger distance corresponded to the scale of the object, indicating that explicit modes may be useful [62]. Also, instead of relying on two-finger touch, we can use D1 and D6 by touching the tablet with the index finger on the dominant hand to fix the axis of rotation, and rotating the tablet with the non-dominant hand to rotate the selected object as if the user is turning an object with a wrench. Overall, participants could understand and transform objects easily (*P*1-5).

Navigate. Navigation was perceived to be a hard task for numerous reasons, except *P*6, “Navigating with 5 fingers is easy and it doesn’t conflict with other tasks.” *P*1 felt that the movement directions were backwards (i.e., that it should have been world-centric, rather than tablet-centric movement, despite the world fading to black), and found it hard to navigate. Moreover, similarly to the transform operation, we let participants rotate, move, and zoom-in/out simultaneously. However, as noted in prior studies, separating DOF could improve the control during navigation [39, 62]. Or using D4 and D9, instead of continuous navigation, a user can select a fixed point on the map shown on the tablet and a tap would instantly teleport the user.

*Menu Navigation. P*1 felt uncomfortable interacting with the back of the tablet due to our custom tracker mounting bar and hook-and-loop fasteners. On the other hand, P6 reported, “[...], sometimes it is hard to hold the tablet in left hand.” Except P4 and P6, participants found menu selection to be difficult. The primary reason for discomfort during the menu selection task was from flipping the tablet and interacting with the the back. Arora et al.[1] speculated about a similar issue in their work. We believe it was cumbersome to hold the tablet with the non-dominant hand, and tracking to interact with the menu was far less reliable than multi-touch on its front. To tackle unreliable tablet tracking and to avoid flipping the tablet, we can use D4 and D6 to select menu items on the front side of the tablet, which allows more precise 2D input and does not rely on 3D position tracking of the tablet with the dominant hand.

Discussion

Overall, the results show participants could use the example TabletInVR interaction vocabulary, as implemented in the proof-of-concept system, to accomplish core 3D modeling operations. This further suggests the associated TabletInVR design space and design dimensions were useful for exploring these new types of interactions.

While our system demonstrated integration of both mid-air hand gestures and tablet input to facilitate 3D solid modelling in VR, user evaluation pointed out some limitations. They can be circumvented using alternative combinations of the design dimensions. For instance, to tackle unreliable tablet 3D position tracking of the tablet, we can use D10 (indirect input) on the tablet screen, which is precise for 2D input (D1) and does not rely on 3D position tracking of the tablet. To tackle issues pertaining to using the dominant hand, we can rearrange roles using D2. To tackle motion sickness, D4 can be used to direct user attention to the device, rather than the surrounding VR environment, and so forth.

As a result, design dimensions would help tackle engineering issues until the technology matures. Moreover, we have presented only a small subset of possible interactions; we believe design dimensions could allow interaction design beyond 3D modeling.

While a comparison of more technically mature and robust TabletinVR systems with the controller is warranted, our work demonstrates the feasibility of using the design dimensions to build a usable interaction vocabulary.

7 CONCLUSION

We are the first to investigate the design of an example interaction vocabulary for using a multi-touch tablet in VR for 3D solid modeling. We approach the design methodically and propose design dimensions that inform the design of our vocabulary, but can also inform the design of alternate vocabularies. We validate this interaction vocabulary with a proof of concept system that addresses the core components of 3D modeling and a user study that shows that the interface is useful in replicating and creating original designs.

Our study also identified some limitations which we discuss with possible solutions, but it also hints at future possibilities in this largely unexplored design space. While our focus was on 3D solid modeling, the design dimensions can also inform vocabularies for other applications like gaming, data visualization, and simulation control. Our work can guide future researchers and designers by extending the VR interaction space beyond traditional input devices.

8 ACKNOWLEDGEMENTS

This work was made possible by NSERC Discovery Grant #2018-05187, the Canada Foundation for Innovation Infrastructure Fund "Facility for Fully Interactive Physio-digital Spaces" (#33151), and Ontario Early Researcher Award #ER16-12-184.

REFERENCES

[1] Rahul Arora, Rubaiat Habib Kazi, Tovi Grossman, George Fitzmaurice, and Karan Singh. 2018. SymbiosisSketch: Combining 2D & 3D Sketching for Designing Detailed 3D Objects in Situ. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 185.

[2] Rahul Arora, Rubaiat Habib Kazi, Fraser Anderson, Tovi Grossman, Karan Singh, and George W Fitzmaurice. 2017. Experimental Evaluation of Sketching on Surfaces in VR. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 5643–5654.

[3] Ilhan Aslan, Ida Buchwald, Philipp Koytek, and Elisabeth André. 2016. Pen+ Mid-Air: An Exploration of Mid-Air Gestures to Complement Pen Input on Tablets. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. ACM, 1.

[4] Rob Aspin and Kien Hoang Le. 2007. Augmenting the CAVE: An initial study into close focused, inward looking, exploration in IPT systems. In *Distributed Simulation and Real-Time Applications, 2007. DS-RT 2007. 11th IEEE International Symposium*. IEEE, 217–224.

[5] Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. 2016. Normaltouch and texturetouch: High-fidelity 3d haptic shape rendering on handheld virtual reality controllers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 717–728.

[6] Hrvoje Benko, Edward W Ishak, and Steven Feiner. 2004. Collaborative mixed reality visualization of an archaeological excavation. In *Proceedings of the 3rd IEEE/ACM international Symposium on Mixed and Augmented Reality*. IEEE Computer Society, 132–140.

[7] Hrvoje Benko, Edward W Ishak, and Steven Feiner. 2005. Cross-dimensional gestural interaction techniques for hybrid immersive environments. In *Virtual Reality, 2005. Proceedings. VR 2005. IEEE*. IEEE, 209–216.

[8] Mark Billinghurst, Sisinio Baldis, Lydia Matheson, and Mark Philips. 1997. 3D palette: a virtual reality content creation tool. In *Proceedings of the ACM symposium on Virtual reality software and technology*. ACM, 155–156.

[9] Mark Billinghurst, Tham Piumsomboon, and Huidong Bai. 2014. Hands in space: Gesture interaction with augmented-reality interfaces. *IEEE computer graphics and applications* 34, 1 (2014), 77–80.

[10] Alexander Bornik, Reinhard Beichel, Ernst Kruijff, Bernhard Reitinger, and Dieter Schmalstieg. 2006. A hybrid user interface for manipulation of volumetric medical data. In *3D User Interfaces, 2006. 3DUI 2006. IEEE Symposium on*. IEEE, 29–36.

[11] Doug A Bowman, Mehdi Setareh, Márcio Serolli Pinho, Ndiwalana Ali, Alex Kalita, Yunha Lee, John Lucas, Matthew Gracey, Malini Kothapalli, Qinwei Zhu, et al. 2003. Virtual-SAP: an immersive tool for visualizing the response of building structures to environmental conditions. In *Virtual Reality, 2003. Proceedings. IEEE*. IEEE, 243–250.

[12] Jeff Butterworth, Andrew Davidson, Stephen Hench, and Marc T Olano. 1992. 3DM: A three dimensional modeler using a head-mounted display. In *Proceedings of the 1992 symposium on Interactive 3D graphics*. ACM, 135–138.

[13] William Buxton and Brad Myers. 1986. A study in two-handed input. In *ACM SIGCHI Bulletin*, Vol. 17. ACM, 321–326.

[14] Li-Chieh Chen, Yun-Maw Cheng, Po-Ying Chu, and Frode Eika Sandnes. 2016. The Common Characteristics of User-Defined and Mid-Air Gestures for Rotating 3D Digital Contents. In *International Conference on Universal Access in Human-Computer Interaction*. Springer, 15–22.

[15] Michael Chen, S Joy Mountford, and Abigail Sellen. 1988. A study in interactive 3-D rotation using 2-D control devices. *ACM SIGGRAPH Computer Graphics* 22, 4 (1988), 121–129.

[16] Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. 2018. CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 654.

[17] Bruno R De Araùjo, Géry Casiez, and Joaquim A Jorge. 2012. Mockup builder: direct 3d modeling on and above the surface in a continuous interaction space. In *Proceedings of Graphics Interface 2012*. Canadian Information Processing Society, 173–180.

[18] Barrett Ens, Juan David Hincapié-Ramos, and Pourang Irani. 2014. Ethereal planes: a design framework for 2D information space in 3D mixed reality environments. In *Proceedings of the 2nd ACM symposium on Spatial user interaction*. ACM, 2–12.

[19] Ajoy S Fernandes and Steven K Feiner. 2016. Combating VR sickness through subtle dynamic field-of-view modification. In *3D User Interfaces (3DUI), 2016 IEEE Symposium on*. IEEE, 201–210.

[20] James D Foley, Victor L Wallace, and Peggy Chan. 1984. The human factors of computer graphics interaction techniques. *IEEE computer Graphics and Applications* 4, 11 (1984), 13–48.

[21] Cédric Foucault, Manfred Micaux, David Bonnet, and Michel Beaudouin-Lafon. 2014. SPad: a bimanual interaction technique for productivity applications on multi-touch tablets. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1879–1884.

[22] Jens Grubert, Lukas Witzani, Eyal Ofek, Michel Pahud, Matthias Kranz, and Per Ola Kristensson. 2018. Text Entry in Immersive Head-Mounted Display-Based Virtual Reality Using Standard Keyboards. *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2018), 159–166.

[23] Yves Guiard. 1987. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of motor behavior* 19, 4 (1987), 486–517.

[24] Genesis Virtual Guide. 2018. Genesis Virtual Guide. (2018). https://www.youtube.com/watch?v=hDxOX8LYKHs.

[25] Chris Harrison, Julia Schwarz, and Scott E Hudson. 2011. TapSense: enhancing finger interaction on touch surfaces. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 627–636.

[26] Otmar Hilliges, Shahram Izadi, Andrew D Wilson, Steve Hodges, Armando Garcia-Mendoza, and Andreas Butz. 2009. Interactions in the air: adding further depth to interactive tabletops. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*. ACM, 139–148.

[27] Ken Hinckley, Michel Pahud, Hrvoje Benko, Pourang Irani, Francois Guimbretiere, Marcel Gavriliu, Xiang'Anthony' Chen, Fabrice Matulic, William Buxton, and Andrew Wilson. 2014. Sensing techniques for tablet+ stylus interaction. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 605–614.

[28] Ken Hinckley, Randy Pausch, Dennis Proffitt, and Neal F Kassell. 1998. Two-handed virtual manipulation. *ACM Transactions on Computer-Human Interaction (TOCHI)* 5, 3 (1998), 260–302.

[29] Ken Hinckley, Koji Yatani, Michel Pahud, Nicole Coddington, Jenny Rodenhouse, Andy Wilson, Hrvoje Benko, and Bill Buxton. 2010. Pen+ touch= new tools. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. ACM, 27–36.

[30] HTC. 2017. "Introducing the Logitech BRIDGE SDK". (2 November 2017). Retrieved Mar 31, 2018 from https://blog.vive.com/us/2017/11/02/introducing-the-logitech-bridge-sdk/.

[31] Hessam Jahani and Manolya Kavakli. 2017. Exploring a user-defined gesture vocabulary for descriptive mid-air interactions. *Cognition, Technology & Work* 20 (2017), 11–22.

[32] Daniel F Keefe, Robert C Zeleznik, and David H Laidlaw. 2007. Drawing on air: Input techniques for controlled 3D line illustration. *IEEE Transactions on Visualization and Computer Graphics* 13, 5 (2007), 1067–1081.

[33] Ji-Sun Kim, Denis Gračanin, Krešimir Matković, and Francis Quek. 2008. Finger Walking in Place (FWIP): A traveling technique in virtual environments. In *International Symposium on Smart Graphics*. Springer, 58–69.

[34] Robert W Lindeman, John L Sibert, and James K Hahn. 1999. Hand-held windows: towards effective 2D interaction in immersive virtual environments. In *Virtual Reality, 1999. Proceedings., IEEE*. IEEE, 205–212.

[35] Robert W Lindeman, John L Sibert, and James K Hahn. 1999. Towards usable VR: an empirical study of user interfaces for immersive virtual environments. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 64–71.

[36] Nicolai Marquardt, Ricardo Jota, Saul Greenberg, and Joaquim A Jorge. 2011. The continuous interaction space: interaction techniques unifying touch and gesture on and above a digital surface. In *IFIP Conference on Human-Computer Interaction*. Springer, 461–476.

[37] John C McClelland, Robert J Teather, and Audrey Girouard. 2017. HaptoBend: shape-changing passive haptic feedback in virtual reality. In *Proceedings of the 5th Symposium on Spatial User Interaction*. ACM, 82–90.

[38] Mark McGill, Daniel Boland, Roderick Murray-Smith, and Stephen Brewster. 2015. A dose of reality: Overcoming usability challenges in vr head-mounted displays. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2143–2152.

[39] Daniel Mendes, Filipe Relvas, Alfredo Ferreira, and Joaquim Jorge. 2016. The benefits of DOF separation in mid-air 3D object manipulation. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*. ACM, 261–268.

[40] Minecraft. 2018. Minecraft Wiki. (2018). https://minecraft.gamepedia.com/Block.

[41] Oculus. 2018. Oculus Locomotion. (2018). https://developer.oculus.com/design/latest/concepts/bp-locomotion.

[42] Thammathip Piumsomboon, Adrian Clark, Mark Billinghurst, and Andy Cockburn. 2013. User-defined gestures for augmented reality. In *IFIP Conference on Human-Computer Interaction*. Springer, 282–299.

[43] Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. 1996. The go-go interaction technique: non-linear mapping for direct manipulation in VR. In *Proceedings of the 9th annual ACM symposium on User interface software and technology*. ACM, 79–80.

[44] Ivan Poupyrev, Numada Tomokazu, and Suzanne Weghorst. 1998. Virtual Notepad: handwriting in immersive VR. In *Virtual Reality Annual International Symposium, 1998. Proceedings., IEEE 1998*. IEEE, 126–132.

[45] Gerhard Reitmayr and Dieter Schmalstieg. 2001. Mobile collaborative augmented reality. In *Augmented Reality, 2001. Proceedings. IEEE and ACM International Symposium on*. IEEE, 114–123.

[46] Élisabeth Rousset, François Bérard, and Michaël Ortega. 2014. Two-finger 3D rotations for novice users: surjective and integral interactions. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces*. ACM, 217–224.

[47] Jaime Ruiz, Yang Li, and Edward Lank. 2011. User-defined motion gestures for mobile interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 197–206.

[48] Markus Sareika and Dieter Schmalstieg. 2010. Bimanual handheld mixed reality interfaces for urban planning. In *Proceedings of the International Conference on Advanced Visual Interfaces*. ACM, 189–196.

[49] Peng Song, Wooi Boon Goh, William Hutama, Chi-Wing Fu, and Xiaopei Liu. 2012. A handle bar metaphor for virtual object manipulation with mid-air interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1297–1306.

[50] Peng Song, Wooi Boon Goh, William Hutama, Chi-Wing Fu, and Xiaopei Liu. 2012. A handle bar metaphor for virtual object manipulation with mid-air interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1297–1306.

[51] Richard Stoakley, Matthew J Conway, and Randy Pausch. 1995. Virtual reality on a WIM: interactive worlds in miniature. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press/Addison-Wesley Publishing Co., 265–272.

[52] Evan Strasnick, Christian Holz, Eyal Ofek, Mike Sinclair, and Hrvoje Benko. 2018. Haptic Links: Bimanual Haptics for Virtual Reality Using Variable Stiffness Actuation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 644.

[53] Hemant Bhaskar Surale, Fabrice Matulic, and Daniel Vogel. 2017. Experimental Analysis of Mode Switching Techniques in Touch-based User Interfaces. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3267–3280.

[54] Ivan E Sutherland. 1965. The ultimate display. *Multimedia: From Wagner to virtual reality* (1965), 506–508.

[55] Zsolt Szalavári and Michael Gervautz. 1997. The Personal Interaction Panel - a Two-Handed Interface for Augmented Reality. *Comput. Graph. Forum* 16 (1997), 335–346.

[56] Anthony Talvas. 2014. *Bimanual haptic interaction with virtual environments*. Ph.D. Dissertation. INSA de Rennes.

[57] HTC VIVE. 2018. HTC Custom Tracker. (2018). https://dl.vive.com/Tracker/Guideline/HTC_Vive_Tracker_Developer_Guidelines_v1.3.pdf.

[58] Julie Wagner, Stéphane Huot, and Wendy Mackay. 2012. BiTouch and BiPad: designing bimanual interaction for hand-held tablets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2317–2326.

[59] Jia Wang and Robert Lindeman. 2014. Coordinated 3D interaction in tablet-and HMD-based hybrid virtual environments. In *Proceedings of the 2nd ACM symposium on Spatial user interaction*. ACM, 70–79.

[60] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. 2018. Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 86.

[61] Jacob O Wobbrock, Meredith Ringel Morris, and Andrew D Wilson. 2009. User-defined gestures for surface computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1083–1092.

[62] Siju Wu, Amine Chellali, Samir Otmane, and Guillaume Moreau. 2015. TouchSketch: a touch-based interface for 3D object manipulation and editing. In *Proceedings of the 21st ACM Symposium on Virtual Reality Software and Technology*. ACM, 59–68.

[63] Yukang Yan, Chun Yu, Xiaojuan Ma, Shuai Huang, Hasan Iqbal, and Yuanchun Shi. 2018. Eyes-Free Target Acquisition in Interaction Space around the Body for Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 42.

[64] Jingjie Zheng, Xiaojun Bi, Kun Li, Yang Li, and Shumin Zhai. 2018. M3 Gesture Menu: Design and Experimental Analyses of Marking Menus for Touchscreen Mobile Interaction. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 249.