


# Experimental Evaluation of Sketching on Surfaces in VR

Rahul Arora<sup>1,2</sup>, Rubaiat Habib Kazi<sup>2</sup>, Fraser Anderson<sup>2</sup>, Tovi Grossman<sup>2</sup>,  
Karan Singh<sup>1</sup>, George Fitzmaurice<sup>2</sup>

<sup>1</sup>DGP Lab, University of Toronto, Toronto, ON, Canada {arorar, karan}@dgp.toronto.edu  
<sup>2</sup>Autodesk Research, Toronto, ON, Canada {firstname.lastname}@autodesk.com



**Figure 1.** Virtual Reality affords the freedom to sketch in unconstrained 3D spaces (left). However, our study indicates that drawing accurately in VR is challenging. Inaccuracies in depth as well as the target planar projection (right) are common.

## ABSTRACT

Sketching in immersive 3D virtual reality (VR) environments has great potential for a variety of interactive 3D design applications. Precisely sketching the intended strokes in mid-air, however, can be a challenge. In this paper, we present a set of controlled studies to analyze the factors affecting human ability to sketch freely in a 3D VR environment. In our first study, we directly compare traditional sketching on a physical surface to sketching in VR, with and without a physical surface to rest the stylus on. Our results indicate that the lack of a physical drawing surface is a major cause of inaccuracies in VR drawing, and that the effect is dependent on the orientation of the drawing surface. In a second experiment, we evaluate the extent to which visual guidance can compensate for the loss of sketching precision in VR. We found that while additional visual guidance improves positional accuracy, it can be detrimental to the aesthetic quality of strokes. We conclude by distilling our experimental findings into design guidelines for sketching tools in immersive 3D environments.

## Author Keywords

Virtual reality; 3D drawing; motor ability; visual factors.

## ACM Classification Keywords

H.5.1 Information interfaces and presentation (e.g., HCI): Artificial, augmented, and virtual realities.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

CHI 2017, May 06–11, 2017, Denver, CO, USA

© 2017 ACM. ISBN 978-1-4503-4655-9/17/05...\$15.00

DOI: <http://dx.doi.org/10.1145/3025453.3025474>


## INTRODUCTION

Recent advancements in Virtual and Augmented Reality (VR/AR) devices have spurred considerable public interest in utilizing the technology for design applications. The availability of consumer-grade devices such as HTC Vive and Microsoft HoloLens has enabled development of mainstream and experimental tools for a variety of visual design tasks such as painting & sketching [43, 44], 3D modeling [44], storytelling [9], and conceptualization [21]. These tools have the potential to change how everyday objects are visualized and designed as they allow users to create, view and modify their designs in 3D at a real-world scale (see Figure 2).

Sketching is a basic task used in many visual design pipelines due to its freeform and expressive nature. Since freehand sketching provides an intuitive method of conceptualizing ideas, there has been considerable research [4, 10, 18, 34, 41, 42] into using freehand sketches to create three-dimensional artifacts by lifting 2D sketches into the third dimension. The HCI community has complemented this research by producing interfaces allowing sketching directly in 3D [15, 25, 36]. Prior works [20, 36] indicate that professional designers are excited by the possibility of using direct 3D input for sketching, but find it difficult, and get frustrated by the lack of control over their strokes (Figure 1). While there have been some quantitative studies exploring how 3D sketching capabilities improve with learning [40] and how force feedback affects precision [27], there has been limited work that evaluates and quantifies the factors which influence drawing in unconstrained, mid-air environments.

In this paper, we perform a set of formal experiments to study the human ability to draw in mid-air VR environments, and explore the impact that physical and visual guidance can have on 3D sketching accuracy. We first conduct a series of

observational sessions with professional designers to understand how professionals approach mid-air sketching. Based on these sessions, we formulate an experiment aimed at quantifying drawing accuracy in VR compared to traditional 2D drawing by varying the presence of a physical drawing surface, in three canonical orientations. We then conduct a second experiment to observe how visual guidance provided by the drawing system in VR impacts drawing accuracy for both planar and non-planar curves, across a greater range of surface orientations.



**Figure 2. VR sketching using consumer-grade tools—Google Tilt Brush (top<sup>1</sup>) and Gravity Sketch (bottom<sup>11</sup>).**

Our first study showed that the mid-air drawing accuracy in VR, measured via the mean overall deviation from a target stroke, decreased by 148% compared to traditional 2D drawing. However, by including a physical drawing surface within a VR environment, the loss of accuracy was reduced to 20%. Our second study found that visual guidance can serve to improve mid-air sketching accuracy, but leads to worse aesthetic quality of strokes, as measured via curve fairness [12]. In both studies, surface orientation was a key factor of the above accuracy levels. Based on these results, we present a broad set of interaction guidelines that can help guide the future design of VR-based design tools.

## RELATED WORK

Our work relates to existing literature on 3D curve creation, design using immersive displays, and drawing ability evaluations. We discuss each of these below.

### 3D Curve Creation

The traditional method for creating curves in 3D is by using a set of interface elements to specify multiple geometric constraints leading to the final 3D curve. Desktop modelling software allows users to drag control points in 3D space to create the desired curvature. Alternative interfaces in the research literature include specialized hardware such as Grossman et al.’s tape based system [16], or using software-based interfaces such as the one proposed by Bae et al. [4] to specify projection planes.

Closer to our topic of study is a class of 3D curve creation techniques that involve freehand 3D motion to specify curves. Due to the lack of 3D display technologies, early

systems, such as the 3-Draw tool [33] were forced to use 2D displays. More recently, 3D display and head tracking technologies have enabled the creation of complex and high quality curve networks [27, 29, 39] using mid-air 3D motion as direct input. Notably, such systems have mostly been explored for generating freeform “organic” curves (Figure 2). We hypothesize that this is because of the innate difference in human drawing ability in 2D compared to 3D, and this forms an important motivation of the initial observation sessions presented in this paper.

## Design using Immersive Display Technologies

Early graphics research produced a number of prototypes for design tools utilizing immersive hardware (displays covering a large field-of-view) [29, 39, 41]. Recent advancements in immersive display and spatial tracking technologies have led to a new spur in commercial, research, as well as hobbyist tools harnessing virtual reality for painting, drawing, and modelling tasks [9, 21, 43–45]. Demand for VR content has inspired the creation of novel VR interfaces for authoring storyboards [17] and animations [46]. However, limited work has investigated the usability and human factors of such 3D design tools. To study designer preference for various UI interfaces, Israel et al. [20] conducted a luminary study involving focus groups of expert designers. While they collected subjective opinions on 3D drawing and sketching, there was no quantitative evaluation of how designers utilized these interaction techniques.

## Drawing Ability Evaluation

There have been many studies on human drawing ability in the fields of motor control, psychology, and HCI. Cohen and Bennett [8] attributed the misperception of the target object as a major reason for drawing inaccuracies. Schmidt et al. [35] focused on expert performance in drawing simple curves in various different projections, and found that even expert artists fail to perceive and/or draw on 2D projections of 3D objects accurately. Fitzmaurice et al. [13] discussed how artists reach optimal orientations in pen-and-paper drawing by rotating the paper, and how digital interfaces for 2D drawing can support this interaction. Our study extends these findings to freehand 3D sketching.

In the VR domain, Keefe et al. [26, 27] studied the impact of force feedback on 3D sketching accuracy by utilizing a Phantom Omni haptic device. In contrast, we compare mid-air sketching to the natural, passive feedback provided by a physical surface, thus better mimicking the traditional drawing setting. Further, unlike a Phantom Omni, this does not limit input range. We also quantify the influence of various factors in isolation, and present novel results on surface-projected sketch accuracy. Jackson and Keefe [22] and Kühnert et al. [28] performed design studies to explore sketching over physical props for rapid prototyping. These works informed the visual guidance in our second study.

Perhaps the investigation closest in application to ours is that of Wiese et al. [40], who studied the learnability of mid-air sketching in a CAVE environment. Our research contributes



new findings on how factors such as physical constraints, visual guidance, orientation, and scale affects mid-air drawing ability in a head-mounted VR environment.

Motor control studies, such as the one by Abend et al. [1], have studied the speed and position profiles of the human hand for various target motions. In the HCI community, various models for understanding speed and accuracy have been proposed for 2D gestures [7, 19], but work in 3D has been limited to low-level motor control evaluations [24].

### INITIAL OBSERVATIONAL SESSIONS

Before conducting our two controlled experiments, we first directly observed artists working within a VR 3D sketching system. The goal of these sessions was to obtain feedback on the challenges and opportunities of 3D freehand sketching. The resulting observations would be used to guide the areas of focus for our quantitative evaluation.

We invited five expert designers (two industrial designers, two concept artists, and an architect) to participate in a design session using Tilt Brush [43] (version 5.4), a VR sketching application. An HTC Vive device was used for the sessions. Participants were asked to generate a 3D sketch of anything related to their domain and expertise in a 60-minute session.

### Observations

In general, the participants were positive about the ability to draw in scale, directly in 3D, utilizing the immersive and freeform nature of VR sketching. Artists drew conceptual models of edge-heavy objects such as cars, an interior design of a room, as well as organic shapes such as shoes and humans. We observed that the drawing characteristics seemed to vary across various positions and orientations of the drawing plane. Participants did mention that accuracy and precision of the strokes in 3D VR environments were more critical, compared to its 2D counterparts:

*“In 3D situations, line accuracy is more important than 2D situations, since you’re conveying more information to the viewer. In 2D, the viewer needs to make a mental leap to go from 2D to 3D. On the contrary, in 3D, because you get more info, you’re looking more precisely at the quality of the line” (P2)*

In general, participants felt that it could be challenging to depict a desired shape, noting in particular that curves that were meant to be planar often ended up as convex (P2, P5). The participants also felt they could not always achieve the intended result due to inaccurate positioning (P1, P2). The designers unanimously agreed that they require more tools for precision and greater controls for meaningful design tasks in VR. Participants suggested having snapping tools (P2, P3), haptic feedback (P1, P4), and projection to existing planes and surfaces (P2, P3, P5). Participants (P2-P5) also reported that ergonomic issues such as neck and shoulder pain may occur when using VR over an extended period.

Our observations are consistent with those that can be made from existing repositories of VR sketching workflows [47]. While artists were able to design successfully, there were certainly struggles when trying to depict a desired curve

accurately in 3D space. In the following sections, we present two controlled experiments to gain a better understanding of the human limits of freehand drawing, which could eventually lead to design recommendations for such systems.

### EXPERIMENT 1: VR VS. TRADITIONAL DRAWING

While the ability to draw non-planar curves makes mid-air sketching unique, planar curves play an integral role for both shape perception and within the 3D design process, as evidenced by numerous previous works [4, 30, 41], as well as our initial observational sessions. In the first experiment, we sought to investigate how mid-air drawing ability of planar curves in VR compares to traditional drawing on a flat surface. In particular, we wanted to study how the presence or absence of a physical drawing surface affects drawing of planar curves. To this end, the participants were exposed to three main experimental conditions. In the *traditional* condition, participants drew on a physical surface. In the *VR* condition, participants used a VR head-mounted display and drew in 3D space. In the *hybrid* condition, participants drew on a physical surface while using the same VR headset.

### Experimental Design

The experiment was designed as a 3 x 3 x 3 x 3 within-subjects study, with independent variables of *Condition* (*traditional*, *hybrid*, *VR*), *Drawing Plane* (*horizontal*, *vertical*, *sideways*), *Shape* (*u-line*, *v-line*, *circle*), and *Size* (*small*: 10cm, *medium*: 30cm, *large*: 60cm). The three drawing plane configurations are illustrated in Figure 3a. All of these planes were located at a comfortable position [6, 32]. For each of these planes, we defined *U-V* coordinate axes that define the *u-line* and *v-line* orientations (Figure 3b).

Participants performed the experiment in a single session lasting 40-50 minutes. The order in which participants were exposed to each *condition* was counterbalanced by a Latin square. Within each condition, the order of the appearance of the drawing planes was again dictated by a Latin square per participant. For each plane orientation, the stroke shape order was randomized. For each stroke shape, participants had to complete three sets of trials, each of which was a random permutation of the three stroke sizes. Overall, each participant drew 243 strokes. For practice, participants drew a medium-sized stroke for each shape, before each condition.


### Participants

In total, 12 able-bodied individuals (8 male, 4 female), aged 22 to 51, participated in the study. All participants were right-handed or ambidextrous, and used their right hand to draw. Participants were 164-183cm tall, and none had sketching experience using VR/AR devices, or professional drawing experience. Participants were paid for participating.


### Apparatus

In the *traditional* condition, participants had access to a large (84”) display, which displayed the target strokes and was used as a physical drawing surface. Drawing was performed using a dry-erase pen augmented with motion-capture markers (Figure 4a). However, strokes were only recorded when the participants pressed a trigger on a HTC Vive

controller (Figure 4b) held in their left hand. In the *hybrid* condition, participants wore the HTC Vive HMD, and the large display was still used as the drawing surface. Simplified 3D models of the display and the tracked pen were rendered via the HMD. In the *VR* condition, no drawing surface was used or displayed, and a matte gray background was rendered. In all the three conditions, the pen position and orientation was tracked at 60 Hz using OptiTrack motion capture cameras. The HTC Vive display had a refresh rate of 90Hz. The software for the experiment was implemented in C# using Unity and SteamVR for rendering and interaction, and OptiTrack's Motive software for motion capture.



**Figure 3.** Alignment of the three drawing planes (a) clockwise from top: vertical, sideways, and horizontal; and arrangement of the three stroke shapes (b) in Experiment 1.



**Figure 4.** Apparatus: The motion-captured pen (a) and the HTC Vive controller (b) used for drawing.

### Procedure

Participants were instructed to stand at a fixed spot before starting to draw in a new drawing plane to ensure controlled results. Without this constraint, participants would be able to change their orientation relative to the drawing surface, thereby confounding the results. For the *horizontal* orientation, the drawing surface was kept at a height of 1m, while for the other two orientations, the surface was arranged so that the center of the stroke was at a height of 1.4m. Following ergonomic guidelines [32], the target strokes were centered approx. 40cm in front of participants.

The trial started with the participant being shown the target stroke along with the starting point, both rendered in black. Participants were instructed to draw all circles in a clockwise direction, while the starting spots on lines were chosen so that the horizontal lines were drawn left-to-right or far-to-near, and the vertical lines top-to-bottom. The target stroke was shown until the participant pressed the trigger in their left hand to initiate drawing. Participants were told to execute the strokes as quickly and accurately as possible. Strokes were rejected automatically if any point on the input was over 20cm away from the target, and the trial was repeated.

For the *traditional* condition, the target strokes were shown on the large screen, and the marker left a visible ink trail. For

the *hybrid* condition, the screen was tracked by motion capture markers, and rendered virtually at the same position in space so that participants actually drew on the physical surface. In the *VR* and *hybrid* conditions, participants used the same pen to draw, but the target and drawn strokes were shown via the HMD.

In the *traditional* condition, an experimenter erased the stroke drawn by the participant after each trial. In the *hybrid* and *VR* conditions, the input stroke was hidden 1 second after the participant finished drawing it.

### Data Preparation

Strokes were recorded as sequences of 3D points sampled at 60Hz. The points were initially transformed to a local coordinate system defined relative to the target stroke. This is defined for lines such that the line starts at  $(0, 0, 0)$  and ends at  $(l, 0, 0)$ , and for circles such that the target stroke is centered at  $(0, 0, 0)$  and the starting point is at  $(-l/2, 0, 0)$ , where  $l$  is the value of the *size* variable. The local XY-plane for the stroke coincides with the drawing plane.

To preprocess the input data before error measures were computed, a median filter with a window size of 6 points (equivalent to 100ms) was applied to filter out any high frequency tracking noise. Some of the strokes in *traditional* and *hybrid* conditions had small tails sticking out of the plane at either ends of the stroke due to slight difference between the participants' intention of beginning/ending drawing and actually pressing/leaving the ON/OFF switch. These tails, found via deviation from mean z-coordinate of the stroke, were removed before further processing. Then, piecewise linear approximation of the stroke length was used to resample all input strokes to 100 equidistant points. Finally, all the strokes were translated such that the starting point matched the starting point of the target stroke. This was done to minimize the impact of positional error caused by participants misjudging the exact position of the target stroke, as well as equipment error such as minor inaccuracies in calibrating the coordinate axes of the motion capture system and the Vive. We denote a resampled and translated stroke  $S$  via the sequence of points  $p_1, p_2, \dots, p_{100}$ , where each point  $p_i = (x, y, z)$  is represented as a 3-tuple.

### Measurements

Repeated Measures-ANOVAs were performed to analyze various error measures. All post-hoc pairwise comparisons were performed using Bonferroni-corrected paired  $t$ -tests. The following measures were computed.

#### Mean Overall Deviation

Mean deviation is defined as the average distance of the (resampled) points of user-drawn stroke from the target stroke. Due to the resampling and positional correction in the pre-processing steps, this measure effectively computes the average distance of the drawn stroke from the target stroke. For a line, this is therefore defined as the average deviation from the local X-axis.

$$\overline{OD}_L = \frac{1}{n} \sum_{i=1}^n \sqrt{(p_i \cdot y)^2 + (p_i \cdot z)^2}$$

For circles, it is the average deviation from the target circle.

$$\overline{OD}_C = \frac{1}{n} \sum_{i=1}^n \sqrt{\left(\sqrt{(p_i \cdot x)^2 + (p_i \cdot y)^2} - l/2\right)^2 + (p_i \cdot z)^2}$$

### Mean Projected Deviation

Since two of the conditions involved a physical constraint that would eliminate any errors in the  $z$  direction, an informative measure of accuracy is also the deviation from the intended shape when the drawn shape is projected onto a plane. This can be computed by setting the  $z$  term in the mean overall deviation equations to zero. Note that for *traditional* and *hybrid* conditions, overall and projected deviations are equivalent, as  $p_i \cdot z$  is always zero.

## Results

### Mean Overall Deviation

The *condition* variable had a significant effect on the mean overall deviation ( $F_{2,22} = 69.6, p < .001$ ). Statistically significant effects were also observed for *drawing plane* ( $F_{2,22} = 8.09, p < .005$ ), *shape* ( $F_{2,22} = 203, p < .001$ ) and *size* ( $F_{2,22} = 169, p < .001$ ). The value for the *VR* condition ( $M = 2.08\text{cm}, SD = 1.71\text{cm}$ ) was much higher than that for the *hybrid* ( $M = 1.01\text{cm}, SD = 1.09\text{cm}$ ) and *traditional* ( $M = 0.84\text{cm}, SD = 0.87\text{cm}$ ) conditions.




Figure 5. Main effects of *condition* (left) and *drawing plane* (right) for mean projected deviation.




Figure 6. *Stroke shape* x *condition* interaction on mean projected deviation.

Post-hoc pairwise comparison showed that all three conditions were significantly different, with *VR* exhibiting over twice the inaccuracy of either of the other two

conditions, showing that moving into the third dimension adds a large amount of inaccuracy. For brevity, we skip further discussion to focus on mean projected deviation.

### Mean Projected Deviation

Significant main effects were observed for all four factors—*condition* ( $F_{2,22} = 23.8, p < .001$ ), *drawing plane* ( $F_{2,22} = 27.5, p < .001$ ), *shape* ( $F_{2,22} = 173, p < .001$ ), and *size* ( $F_{2,22} = 133, p < .001$ ). Interaction effects were observed for both *condition* x *shape* ( $F_{4,44} = 6.89, p < .001$ ), *condition* x *size* ( $F_{4,44} = 9.60, p < .001$ ), but not *condition* x *plane* ( $F_{4,44} = 1.43, p = .24$ ).

The *VR* condition resulted in the highest deviation ( $M = 1.29\text{cm}, SD = 1.32\text{cm}$ ) followed by *hybrid* ( $M = 1.01\text{cm}, SD = 1.09\text{cm}$ , and *traditional* ( $M = 0.84\text{cm}, SD = 0.87\text{cm}$ ). Post-hoc pairwise comparison showed significant differences between all *conditions* ( $p < .05$ ) (Figure 5, left).

Stated differently, there was a 20.2% decrease in accuracy by wearing an HMD and rendering the environment virtually, and a further 27.2% decrease by removing the supporting physical surface. These results demonstrate that the difficulty of in-air 3D sketching is two-fold, impacted by both motor and visual limitations. Visually, the user no longer has direct line of sight to their hand and pen, and must rely on virtual depth perception cues. From a motor standpoint, the lack of a physical surface introduces even further inaccuracies.

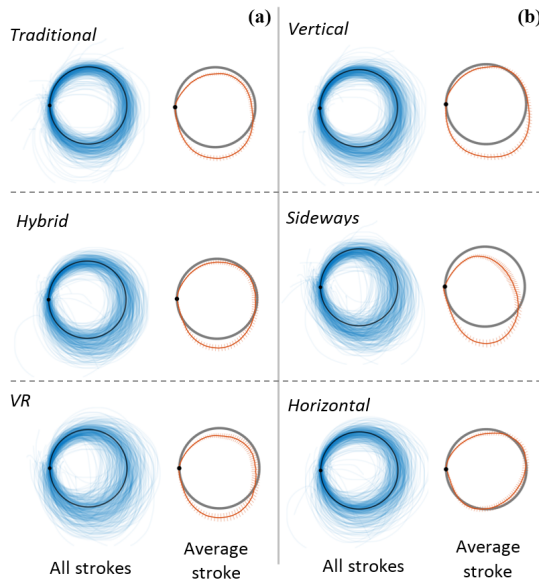
With respect to orientation, the *sideways* drawing plane had the largest mean projected deviation ( $M = 1.30\text{cm}, SD = 1.30\text{cm}$ ), and was significantly different ( $p < .01$ ) from both the *horizontal* ( $M = 0.95\text{cm}, SD = 0.94\text{cm}$ ) and *vertical* ( $M = 0.88\text{cm}, SD = 0.98\text{cm}$ ) planes which were not significantly different from each other (see Figure 5, right). The irregular movements required to draw on the *sideways* drawing plane likely contributed to these results.

Figure 6 shows the effects of *shape* for each condition. As one may expect, the *circle* caused the most difficulty, ( $M = 1.90\text{cm}, SD = 1.38\text{cm}$ ) as compared to one-dimensional lines (combined  $M = 0.62\text{cm}, SD = 0.62\text{cm}$ ). It can be seen that the effect is amplified in the *VR* condition, as indicated by the significant *condition* x *size* interaction. All values of *Size* were significantly different ( $p < .001$ ), with the means following the order *small*: 0.54cm ( $SD = 0.51\text{cm}$ ), *medium*: 0.94cm ( $SD = 0.88\text{cm}$ ), and *large*: 1.65cm ( $SD = 1.46\text{cm}$ ). This indicates that drawing a stroke of larger size is inherently more difficult, hinting at a trend similar to that observed by Cao and Zhai [7] for 2D gestures of much-smaller sizes. The trend is probably due to the increasing effect of the stroke being dictated by the natural arcs of arm movement.

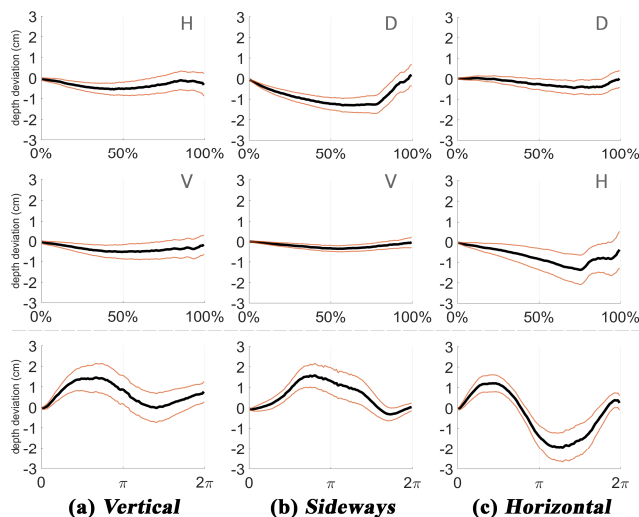
### Shape Trends

Beside aggregated measures of accuracy, it is also interesting to observe trends in deviation from target as participants progressed through the strokes. We visualize mean projected

deviation trends for circles in Figure 7a. A noticeable effect is the consistency of the average circle across the *conditions*. In contrast, the average circles for the three *drawing planes* are visibly different (Figure 7b). This may be due to different muscle groups needed to draw on each of the planes.



**Figure 7. Visualization of projected deviation for circles.** Circles drawn in different *conditions* (a), and *drawing planes* (b). For each row, all circles (left), and average circle with 95% confidence interval (right) are shown.



**Figure 8. Depth deviation trends for different drawing planes and shapes.** Row 1-2: lines, row 3: circles. Key for lines—H: horizontal left-to-right, V: vertical, and D: horizontal (depth) far-to-near.

For the *VR* condition, we can also look at the deviation in depth across the stroke progress. Figure 8 illustrates this effect of drift for each orientation and shape. In particular, it can be seen that the same lines drawn in different imaginary planes have very different characteristics. This may be an


effect of pen grip, which was informally observed to change with the drawing planes, even in the *VR* condition.

## EXPERIMENT 2: FACTORS OF VR DRAWING

Our first experiment showed that there are both motor and visual challenges involved with drawing in VR. In this second experiment, we focus on the VR drawing condition in greater depth, and determine if some of the challenges can be alleviated with enhanced visual guidance. Expanding on Experiment 1, we also examine a larger variety of plane orientations and the effect of drawing non-planar strokes.

### Visual Guidance

A central challenge of working in 3D is being provided with adequate depth cues and visual feedback. We therefore explored combinations of two forms of visual guidance resulting in four visual guidance conditions (Figure 9).



**Figure 9. Visual guidance variants tested in Experiment 2.** Left to right: *none*, *surface*, *stroke*, and *both*.

### Surface Grid

Previous work in sketch-based modelling uses rendered surfaces to help users anchor their strokes in 3D [4]. We sought to examine the effect of visual guidance plane on the characteristics of sketched strokes. We use a grid line pattern as the rendered surface, as prior research [2, 5] suggests it as an effective texture for the perception of a surface.

### Target Stroke

A more detailed form of visual guidance is to render and trace over the actual target stroke. In VR, this condition simulates drawing over virtual objects, inspired by existing literature [22]. Examples of this task include marking out boundaries of CT scan [23], or using an imported image or model as scaffolding to guide freehand strokes [15].

In addition to these two guidance conditions, we also test a condition that combined them (*both*), and a condition that provided no visual guidance (*none*).

### Drawing Surface Orientations

To gain a deeper understanding of how the orientation of the drawing surface affects accuracy, we tested 13 different surface orientations. We positioned the drawing surfaces tangential to a hemisphere of radius 40cm, placed approx. 15cm in front of the participant's head. The orientation is defined using the pair of angles  $(\phi, \theta)$  by which the plane was rotated around the center of the hemisphere along the vertical (down-to-up) and horizontal (left-to-right) axes, respectively. Both  $\phi$  and  $\theta$  were defined to be zero for the vertical plane in front of the participant. The various orientations studied were (Figure 10a):



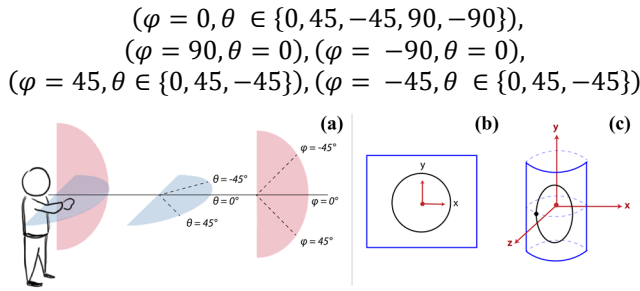


Figure 10. Orientations (a) and surface shapes—*flat* (b) and *curved* (c)—used in the Experiment 2.

### Experimental Design

The experiment was designed as a 4x13x2 within-subject study. The independent variables were *visual guidance* (*none*, *surface*, *stroke*, *both*), *surface orientation* (see above) and *surface shape* (*curved* or *flat*). The curved surface consisted of a cylindrical surface of radius  $R_0 = 15\text{cm}$  (Figure 10c). The target shape was fixed as a 30cm circle.

The experiment was divided into four blocks, one for each of the *visual guidance* conditions. The ordering of *visual guidance* was counterbalanced using a balanced Latin square. For each level of visual guidance, the participants were exposed to 13 orientations of the drawing surface, ordered randomly. For each of these orientations, they had to complete three sets of trials, each of which involved drawing two circular strokes—one on a planar surface (*flat*), and the other on the cylindrical surface (*curved*). The surface shapes were ordered randomly. Overall, each participant executed 312 strokes, in a single session lasting 50-80 minutes.

### Participants

We recruited 12 able-bodied participants (8 male, 4 female), aged 22 to 53. None of the participants had experience using VR/AR for sketching, or professional experience sketching or painting. Participants were paid for their participation.

### Procedure

Participants were instructed to stand at a fixed spot for the duration of the experiment, and to avoid moving their lower body as much as possible. The procedure was otherwise the same as the first experiment: participants used a trigger switch in their left hand to engage the pen, and the drawn stroke was shown to them in real-time. Each trial consisted of drawing a circle of diameter 30cm as a single stroke. The apparatus was the same as the first experiment, using an HTC Vive and OptiTrack system for display and tracking.

### Data Preparation

The filtering procedure, as discussed in Experiment 1, was performed to remove tracking noise from the data. However, strokes were not translated to align their starting points, to observe how the presence or absence of certain visual guidance affects depth perception and positioning accuracy.

For comparing stroke characteristics for the two values of the *surface shape* variable, the circles drawn on the *curved* (cylindrical) surfaces were transformed to a planar circle by

unwrapping the surface [3] such that it becomes the XY plane in the new coordinate system, while the Z direction remains the same. This unwrapping transforms curves drawn in the *curved* condition into an equivalent coordinates system as the curves drawn in the *flat* condition.

### Measurements

The main measurement is mean projected deviation, as defined for Experiment 1, which represent the components of the overall deviation in the local XY plane. We also look at depth deviation (defined below), which represents the overall deviation in the local Z direction. We do not discuss overall deviation, as its trends were very similar to that of mean projected deviation. In addition to these accuracy levels, we also examine the aesthetic quality of the drawn strokes, and the stroke execution time.

#### Mean Depth Deviation

The mean depth deviation refers to mean deviation of the stroke from the target in local Z-direction ( $1/n * \sum_{i=1}^n |p_i.z|$ ). This measure will help estimate the effectiveness of visual guidance to aid depth perception during the drawing task.

#### Mean Fairness Deviation

Fairness of a curve is an important measure of its aesthetic quality [12, 31], useful in many design applications. The fairness of a curve is defined using the smoothness of its curvature. We use a very simple approximation to compute this. Formally, if the curvature of the input curve at point  $p_j$  is  $C_j$ , then we define mean fairness deviation as:

$$\overline{FD} = \left(\frac{1}{n-1}\right) \sum_{i=1}^{n-1} \frac{|C_{i+1} - C_i|}{\text{avg}(C_i, C_{i+1})}.$$

Intuitively, this measure penalizes frequent curvature changes in the input stroke, which are undesirable in design applications. A perfect circle has a fairness deviation of 0.

#### Stroke Execution Time

Stroke execution time is the total time spent in executing a stroke. Measuring this allowed us to estimate the extent to which strokes were performed as rapid arm movements in contrast to participants correcting themselves mid-stroke.

### Results

#### Mean Projected Deviation

Significant main effects were observed for *surface shape* ( $F_{1,11} = 59.3, p < .001$ ), *visual guidance* ( $F_{3,33} = 81.8, p < .001$ ), as well as for *surface orientation* ( $F_{12,132} = 2.74, p < .01$ ). The *angle x surface shape* interaction was also significant ( $F_{12,132} = 3.24, p < .001$ ).

For *visual guidance* conditions, the greatest error was with *none* ( $M = 2.02\text{cm}, SD = 1.05\text{cm}$ ), followed by *surface* ( $M = 1.67\text{cm}, SD = 0.72\text{cm}$ ), *stroke* ( $M = 0.87\text{cm}, SD = 0.51\text{cm}$ ), and *both* ( $M = 0.84\text{cm}, SD = 0.52\text{cm}$ ) (Figure 11, left). Post-hoc comparisons revealed that all pairs other than (*stroke*, *both*) were significantly different. This suggests that the drawing surface provides useful guidance for participants, and showing the target stroke helps even more.

However, showing the drawing surface when the target curve is already visible may not lead to additional accuracy.

For *surface shape*, mean projected deviation was higher for *curved* ( $M = 1.63\text{cm}, SD = 0.97\text{cm}$ ), compared to *flat* ( $M = 1.07\text{cm}, SD = 0.71\text{cm}$ ). This may be due to additional mental and motor overhead costs associated with drawing a non-planar stroke.

In contrast to the large effect of *drawing plane* observed in Experiment 1, the difference between various *surface orientations* was significant, but smaller. The maximum difference between any pair of surface orientation values was 0.32cm. This may be because the surfaces in this experiment were positioned to have projections of comparable sizes on the participants' eyes. Consistent with the first experiment, accuracy was greatest at the horizontal plane.

#### Mean Depth Deviation

The mean depth deviation was also significantly affected by *visual guidance* ( $F_{3,33} = 23.8, p < .001$ ), *surface orientation* ( $F_{12,132} = 4.52, p < .001$ ), and *surface shape* ( $F_{1,11} = 8.75, p < .05$ ). Post-hoc comparisons revealed that *none* was significantly less accurate than each of the three visual guidance conditions. However, the three visual guidance conditions were not significantly different from one another (Figure 11, center). The exact values followed the order: *none* ( $M = 2.11\text{cm}, SD = 1.02\text{cm}$ ), *surface* ( $M = 1.37\text{cm}, SD = 1.11\text{cm}$ ), *stroke* ( $M = 1.01\text{cm}, SD = 0.58\text{cm}$ ), and *both* ( $M = 0.91\text{cm}, SD = 0.46\text{cm}$ ).

For *surface shape*, the mean depth deviation for *curved* surfaces ( $M = 1.27\text{cm}, SD = 0.86\text{cm}$ ) was actually lower than for *planar* surfaces ( $M = 1.43\text{cm}, SD = 1.04\text{cm}$ ). However, there was a significant interaction effect observed between *surface shape* and *visual guidance* ( $F_{3,33} = 4.93, p < .01$ ). *Curved* surfaces exhibited higher depth deviation than *flat* ones when no visual guidance was provided, but lower depth deviation when any additional visual guidance was provided. This shows that with adequate

visual feedback, users may be able to draw non-planar curves as accurately as planar ones.

The mean depth deviation trend for surface orientation is visualized in Figure 11, right. This data shows that sketching in the vertical plane is least accurate ( $M = 1.66\text{cm}, SD = 1.35\text{cm}$ ), and accuracy gradually increases as the orientation moves away from the vertical plane in all directions. Consistent with the first study, accuracy is greatest at the horizontal plane ( $M = 1.02\text{cm}, SD = 0.63\text{cm}$ ).

#### Mean Fairness Deviation

Both *visual guidance* ( $F_{3,33} = 20.6, p < .001$ ) and *surface shape* ( $F_{1,11} = 22.6, p < .001$ ) had a significant effect on Mean Fairness Deviation. Drawn curves were significantly fairer when the target curve was invisible when drawing ( $M_{\text{none}} = 0.29$ ;  $M_{\text{surface}} = 0.32$ ), as compared to when it was visible ( $M_{\text{stroke}} = 0.37$ ;  $M_{\text{both}} = 0.38$ ). This followed our informal observation, as participants focused on staying close to the target curve when that guide was displayed.

For surface shapes, we observed that circles drawn in *planar* condition ( $M = 0.32, SD = 0.10$ ) were fairer than those on *curved* surfaces ( $M = 0.36, SD = 0.09$ ). Lastly, *surface orientation* was found to be significant as well ( $F_{11,132} = 7.73, p < .001$ ). In particular, drawing on the vertical plane ( $(\phi, \theta) = (0, 0)$ ) produced the fairest curves.

#### Stroke Execution Time

Total time spent executing curves on the two *surface shapes* was significantly different ( $F_{1,11} = 34.3, p < .001$ ). This is somewhat expected since the planar circles can be completed with a simple movement, while non-planar strokes on the *curved* surface require more careful manipulation of the arm. The level of *visual guidance* also significantly influenced ( $F_{1,11} = 20.2, p < .001$ ) the curve execution time, with participants spending more time with increase in visual guidance ( $M_{\text{none}} = 3.31\text{s}$ ,  $M_{\text{surface}} = 4.11\text{s}$ ,  $M_{\text{stroke}} = 5.58\text{s}$ ,  $M_{\text{both}} = 6.02\text{s}$ ). All but the last two differences were statistically significant.


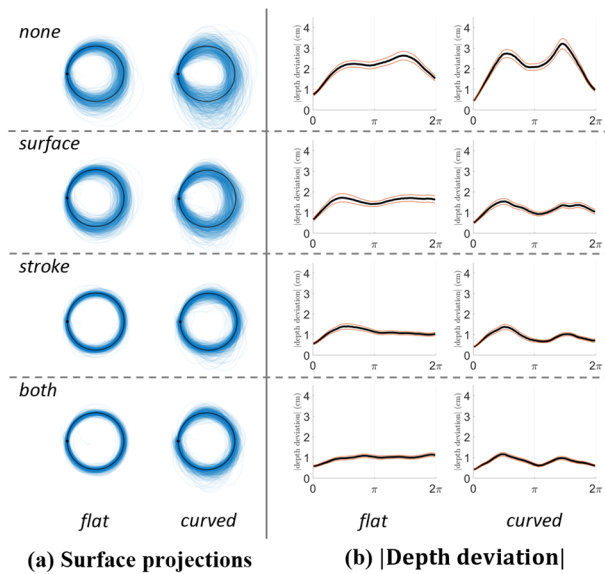


Figure 11. Main effects of *visual guidance* on mean projected deviation (left) and mean fairness deviation (center), and of *surface orientation* on mean depth deviation. The observed trend for *visual guidance* was exactly the opposite for projection and fairness deviations. Depth deviation improved when moving away from the vertical plane in front of the participants.



**Figure 12. Effect of visual guidance and surface shape. Visualization of surface projections of all strokes (a), and trends of depth deviation with 95% confidence intervals (b). Notice the interaction effect between the two variables on the depth deviation measure.**

While curve tracing was slower than drawing without any guidance, similar to Viviani and Terzuolo's [38] observation for curves on a plane, the magnitude of the difference was much smaller than the tenfold difference they found.

Stroke execution time correlated strongly with mean fairness deviation ( $r = .82, p < .001$ ), but negatively correlated weakly with mean overall deviation ( $r = -.38, p < .001$ ). Contrary to the observed trend [7] for gestures, this indicates that spending more time on the stroke leads to marginally better accuracy. However, it leads to a high chance of poor curve quality. This is likely due to the difficulty in keeping a stylus stable during slow movements in 3D space.

#### Visual Trends

Similar to the first experiment, visualization of data reveals interesting trends in the drawing characteristics. We visualize the surface projections of all drawn circles in Figure 12a. In Figure 12b, we illustrate the associated depth deviations. This visualization reveals an interesting trend with respect to depth deviation: the maximum depth deviation lies around  $\pi/2$  radians from the start, and the minimum occurs around  $\pi$  radians. For *curved* surfaces, another local maximum exists around  $3\pi/2$  radians. A similar, albeit more noisy, trend of two characteristic jumps at  $\pi/2$  and  $3\pi/2$  radians was also seen for curvature. This leads to a hypothesis that participants followed an overly linear path for a quarter of the stroke before making a sudden direction change and continuing on towards linearity again, and repeated this process for the next half as well.

#### DISCUSSION AND FUTURE WORK

The initial pilot observation clearly points out that, while designers appreciate the freeform nature of 3D sketching in

VR, precision and control is often required for meaningful design tasks. Designers expressed interest in switching back and forth between freeform and controlled sketching modes.

In our experiments, we observed that the lack of a physical drawing plane, surface position, shape and orientation, stroke size, and visual guidance were important factors affecting drawing ability in VR. While a physical drawing surface improves accuracy by 22% as compared to unguided mid-air drawing, a virtual surface itself, easily incorporated in many design applications, can improve accuracy by as much as 17%. In applications where the target curve is known a priori, showing that curve can further boost accuracy, but may lead to aesthetically poor strokes. It should be noted that our results characterize strokes drawn by amateurs with no prior VR-sketching experience. It is possible that mid-air drawing precision improves with practice and experience. Wiese et al. [40] reported encouraging results on learnability of freehand 3D sketching in a short-term study conducted on design students. Future work on VR sketching learnability can gain from our experience by benchmarking against our results.

While primarily targeted towards sketching applications, our results for conditions with no visual guidance could also be utilized for estimating user precision in performing 3D gestures in VR. Thus, these results could apply to a wide variety of use-cases, such as mode switches or command invocations, or navigation gestures. These use cases result in potential applicability of the results to a number of VR domains such as motor skill training [14], animation and modeling [44, 46], and data visualization [11].


#### Physical Surface

During our observational sessions, designers recommended projecting their strokes into desired virtual planes for greater accuracy. However, our experimental study indicates that, compared to accuracy achieved in traditional sketching, projecting to virtual planes in VR is 53% worse, while drawing on a physical surface in VR is only 20% worse. This is a strong motivation for devices combining existing drawing methods in VR. One potential approach is to hold a clipboard, or other rigid surface in the non-dominant hand to rest upon [33]. An alternative approach is dynamically configurable drawing surfaces actuated by robotic arms (Figure 13a). While the former potentially allows for higher mobility, the latter may be preferable in long design sessions where fatigue becomes a major factor. Such surfaces could potentially provide a more natural haptic feedback without limiting input range, unlike existing active haptic devices used to aid 3D sketching [27].

#### Visual Guidance

Our experimental results indicate that visual guides, such as grids and scaffolding curves, can help to position strokes more accurately, increasing accuracy by 17% and 57%, respectively. Visual guidance may be augmented by snapping techniques, such as the recently implemented "Guides" feature in Tilt Brush [43], which snaps strokes to nearby surface widgets. However, such guides should be

used with caution. In particular, tracing over an existing curve adversely affects the drawing quality. Rendering a grid may provide the right balance between accuracy and aesthetics. These results are in contrast to assumptions for 2D sketching [37] for which curve smoothing algorithms are based. Therefore, novel stroke processing methods may be needed to account for the lack of curve fairness.



**Figure 13: Examples of VR interface elements based on our design suggestions: use of an actuated physical support (a), automatic surfacing with semi-transparent rendering (b), and positioning strokes for maximum projection size (c).**

#### Depth Perception

Depth perception is critical for VR design applications. Our hybrid condition from the first experiment validated that the difficulties of drawing in VR are not solely due to the lack of physical constraint, but also due to the visual impairments. For example, we observed that participants were not able to match the endpoint of their drawn circles to the starting point. Participants in our initial observation sessions had similar complaints about depth perception, more so when the scene had many curves. Advanced visual guides, beyond the approaches tested in our second experiment, could be explored. A common method employed in desktop-based modelling applications is fog rendering, to blur distant objects. In an immersive environment, depth may also be effectively conveyed via a 3D grid, which cannot be employed in 2D due to the resulting visual clutter. While drawing curve networks or concept sketches, automatic surfacing with semi-transparent rendering may also be useful for indicating relative depth (see Figure 13b).

#### Orientation, Position, and Navigation

The drawing plane orientation had a significant effect on accuracy in both experiments. In general, the familiar horizontal orientation was most accurate, with vertical orientations performing worse. The effect was most prominent in the first study when participants had to draw on the sideways plane that was centered with their body. In contrast, the planes in experiment 2 had approximately the same projection on the user's eyes, and all of them were at a comfortable motor distance. Bae et al. [4] defined the "sketchability" measure for 2D sketches, which states that a good view for sketching 2D strokes is one where the sketch surface has a large visible projection. The trend we observed hints that the same model applies for sketching in an immersive VR environment as well. Thus, it would be interesting to explore how VR design tools could equip designers with simple navigation tools that enable them to snap or project the drawing planes to positions and orientations with higher "sketchability" (Figure 13c).

#### Drawing Scale

Our experiments provide evidence that drawing in large scale leads to a sharp increase in drawing inaccuracies. This can be partly attributed to the tendency of the human arm to follow a natural arc. This might lead to the conclusion that, contrary to prior recommendations [20, 36], full-scale design in VR is not a good strategy, if precision is desired. Instead, we suggest that VR design tools allow users to draw in a small manageable scale to improve accuracy and comfort, while allowing users to quickly switch to full-scale view to utilize the visualization benefits afforded by VR [20, 36].

#### Non-planar Strokes

While a large number of design tasks can be performed using planar strokes only [10, 41], previous work suggests that non-planar strokes occur commonly in scenarios such as automobile design [15] and organic shapes [21]. Since our experiments indicate that drawing on planar surfaces produces more accurate and fair curves, we believe that mid-air design tools could often incorporate methods to allow planar curves to be projected onto curved surfaces. However, our results also showed that with adequate visual feedback, users may be able to accurately draw non-planar curves in mid-air. Previous work utilizing 2D displays have used orthographic projections of surfaces onto planes [15], or level set representations of 3D spaces around a surface [34]. How well these techniques relate to artist intent in VR is an interesting avenue to explore in the future. Finally, we only studied a fundamental set of non-planar strokes—circles projected onto a constant curvature convex surface. While the impact of factors studied here is likely independent of shape complexity, more complex 3D strokes are an interesting topic for future investigations. For instance, it is unknown how surface curvature and torsion affect stroke performance.

#### CONCLUSION

By taking away the flatness of paper, VR sketching tools enable artists and designers to dive in and create in free space without any constraints. While this offers immense freedom and expressiveness, our observations indicate that sketching precise curves for design tasks in VR is very challenging. Our paper quantitatively explores multiple factors affecting the imprecision in 3D sketching in VR, including scale, lack of physical surface, orientation, and planarity. We hope that this understanding would lay out the foundation for building tools to equip designers with the desired precision and controls. Based on our analysis, we also discuss a number of design and interaction guidelines that could potentially alleviate the precision problem for sketching in VR and warrant further investigation. However, further studies are required to understand the effectiveness of these solutions to respect artist intent and maintain fidelity. We hope that our work would take this medium forward by making it more receptive to the needs of creative design.

#### ACKNOWLEDGEMENTS

R.A. was funded in part by MITACS Accelerate IT07655.



## REFERENCES

1. W. Abend, E. Bizzi, and P. Morasso. 1982. Human Arm Trajectory Formation. *Brain* 105, 2: 331–348. <http://doi.org/10.1093/brain/105.2.331>
2. George J. Andersen, Myron L. Braunstein, and Asad Saidpour. 1998. The perception of depth and slant from texture in three-dimensional scenes. *Perception* 27, 9: 1087–1106.
3. Tom M. Apostol and Mamikon A. Mnatsakanian. 2007. Unwrapping curves from cylinders and cones. *American Mathematical Monthly* 114, 5: 388–416.
4. Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2008. I Love Sketch: as-natural-as-possible sketching system for creating 3d curve models. *Proceedings of the 21st annual ACM symposium on User interface software and technology*, ACM, 151–160. <http://doi.org/10.1145/1449715.1449740>
5. A. Bair and D. House. 2007. Grid With a View: Optimal Texturing for Perception of Layered Surface Shape. *IEEE Transactions on Visualization and Computer Graphics* 13, 6: 1656–1663. <http://doi.org/10.1109/TVCG.2007.70559>
6. Barbara Bradley. 2003. *Drawing People*. North Light Books, Cincinnati, Ohio.
7. Xiang Cao and Shumin Zhai. 2007. Modeling Human Performance of Pen Stroke Gestures. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, ACM, 1495–1504. <http://doi.org/10.1145/1240624.1240850>
8. Dale J. Cohen and Susan Bennett. 1997. Why can't most people draw what they see? *Journal of Experimental Psychology: Human Perception and Performance* 23, 3: 609.
9. Josh Constine. Oculus “Quill” Turns VR Painting Into Performance Art. *TechCrunch*. Retrieved September 20, 2016 from <http://social.techcrunch.com/2016/01/26/oculus-quill/>
10. Chris De Paoli and Karan Singh. 2015. SecondSkin: Sketch-based Construction of Layered 3D Models. *ACM Trans. Graph.* 34, 4: 126:1–126:10. <http://doi.org/10.1145/2766948>
11. C. Donalek, S. G. Djorgovski, A. Cioc, et al. 2014. Immersive and collaborative data visualization using virtual reality platforms. *2014 IEEE International Conference on Big Data (Big Data)*, 609–614. <http://doi.org/10.1109/BigData.2014.7004282>
12. Gerald Farin and Nickolas Sapidis. 1989. Curvature and the Fairness of Curves and Surfaces. *IEEE Comput. Graph. Appl.* 9, 2: 52–57.
13. George W. Fitzmaurice, Ravin Balakrishnan, Gordon Kurtenbach, and Bill Buxton. 1999. An Exploration into Supporting Artwork Orientation in the User Interface. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*, ACM, 167–174. <http://doi.org/10.1145/302979.303033>
14. Anthony G. Gallagher, E. Matt Ritter, Howard Champin, et al. 2005. Virtual Reality Simulation for the Operating Room: Proficiency-Based Training as a Paradigm Shift in Surgical Skills Training. *Annals of Surgery* 241, 2: 364–372.
15. Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. 2002. Creating Principal 3D Curves with Digital Tape Drawing. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*, ACM, 121–128. <http://doi.org/10.1145/503376.503398>
16. Tovi Grossman, Ravin Balakrishnan, and Karan Singh. 2003. An Interface for Creating and Manipulating Curves Using a High Degree-of-freedom Curve Input Device. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*, ACM, 185–192. <http://doi.org/10.1145/642611.642645>
17. Rorik Henrikson, Bruno De Araujo, Fanny Chevalier, Karan Singh, and Ravin Balakrishnan. 2016. Multi-device storyboards for cinematic narratives in VR. *In proceedings of UIST 2016*.
18. Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 1999. Teddy: A Sketching Interface for 3D Freeform Design. *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*, ACM Press/Addison-Wesley Publishing Co., 409–416. <http://doi.org/10.1145/311535.311602>
19. Poika Isokoski. 2001. Model for Unistroke Writing Time. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*, ACM, 357–364. <http://doi.org/10.1145/365024.365299>
20. J.H. Israel, E. Wiese, M. Mateescu, C. Zöllner, and R. Stark. 2009. Investigating three-dimensional sketching for early conceptual design—Results from expert discussions and user studies. *Computers & Graphics* 33, 4: 462–473. <http://doi.org/10.1016/j.cag.2009.05.005>
21. B. Jackson and D. F. Keefe. 2016. Lift-Off: Using Reference Imagery and Freehand Sketching to Create 3D Models in VR. *IEEE Transactions on Visualization and Computer Graphics* 22, 4: 1442–1451. <http://doi.org/10.1109/TVCG.2016.2518099>
22. Bret Jackson and Daniel F. Keefe. 2011. Sketching Over Props: Understanding and Interpreting 3D Sketch Input Relative to Rapid Prototype Props. *IUI 2011 Sketch Recognition Workshop*.
23. Peter Karasev, Ivan Kolesov, Karl Fritscher, Patricio Vela, Phillip Mitchell, and Allen Tannenbaum. 2013. Interactive medical image segmentation using PDE control of active contours. *IEEE transactions on medical imaging* 32, 11: 2127–2139.
24. Raghavendra S. Kattinakere, Tovi Grossman, and Sriram Subramanian. 2007. Modeling steering within above-the-surface interaction layers. *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM, 317–326. <http://doi.org/10.1145/1240624.1240678>

25. Daniel F. Keefe, Daniel Acevedo Feliz, Tomer Moscovich, David H. Laidlaw, and Joseph J. LaViola Jr. 2001. CavePainting: a fully immersive 3D artistic medium and interactive experience. *Proceedings of the 2001 symposium on Interactive 3D graphics*, ACM, 85–93. <http://doi.org/0.1145/364338.364370>
26. Daniel F. Keefe and David H. Laidlaw. 2007. Poster: Analysis of Performance in Precise 3D Curve Input Tasks in Virtual Reality. *IEEE Visualization (Best Posters Session)*.
27. Daniel Keefe, Robert Zeleznik, and David Laidlaw. 2007. Drawing on air: Input techniques for controlled 3D line illustration. *IEEE Transactions on Visualization and Computer Graphics* 13, 5: 1067–1081.
28. T. Kuehnert, S. Rusdorf, and G. Brunnett. 2011. Virtual Prototyping of Shoes. *IEEE Computer Graphics and Applications* 31, 5: 30–42. <http://doi.org/10.1109/MCG.2010.81>
29. Jung-hoon Kwon, Han-wool Choi, Jeong-in Lee, and Young-Ho Chai. 2005. Free-Hand Stroke Based NURBS Surface for Sketching and Deforming 3D Contents. In *Advances in Multimedia Information Processing - PCM 2005*, Yo-Sung Ho and Hyoung Joong Kim (eds.). Springer Berlin Heidelberg, 315–326.
30. James McCrae, Nobuyuki Umetani, and Karan Singh. 2014. FlatFitFab: Interactive Modeling with Planar Sections. *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*, ACM, 13–22. <http://doi.org/10.1145/2642918.2647388>
31. J. McCrae and K. Singh. 2008. Sketching piecewise clothoid curves. *Proceedings of the Fifth Eurographics conference on Sketch-Based Interfaces and Modeling*, Eurographics Association, 1–8. <http://doi.org/10.2312/SBM/SBM08/001-008>
32. Julius Panero and Martin Zelnik. 1979. *Human Dimension and Interior Space: A Source Book of Design Reference Standards*. Watson-Guptill, New York.
33. Emanuel Sachs, Andrew Roberts, and David Stoops. 1991. 3-Draw: A Tool for Designing 3D Shapes. *IEEE Comput. Graph. Appl.* 11, 6: 18–26. <http://doi.org/10.1109/38.103389>
34. Johannes Schmid, Martin Sebastian Senn, Markus Gross, and Robert W. Sumner. 2011. OverCoat: An Implicit Canvas for 3D Painting. *ACM SIGGRAPH 2011 Papers (SIGGRAPH '11)*, ACM, 28:1–28:10. <http://doi.org/10.1145/1964921.1964923>
35. Ryan Schmidt, Azam Khan, Gord Kurtenbach, and Karan Singh. 2009. On Expert Performance in 3D Curve-drawing Tasks. *Proceedings of the 6th Eurographics Symposium on Sketch-Based Interfaces and Modeling (SBIM '09)*, ACM, 133–140. <http://doi.org/10.1145/1572741.1572765>
36. Shun'ichi Tano, Shinya Yamamoto, Junko Ichino, Tomonori Hashiyama, and Mitsuru Iwata. 2013. Truly Useful 3D Drawing System for Professional Designer by “Life-Sized and Operable” Feature and New Interaction. In *Human-Computer Interaction – INTERACT 2013*. Springer Berlin Heidelberg, 37–55.
37. Yannick Thiel, Karan Singh, and Ravin Balakrishnan. 2011. Elasticurves: exploiting stroke dynamics and inertia for the real-time neatening of sketched 2D curves. *Proceedings of the 24th annual ACM symposium on User interface software and technology*, ACM, 383–392. <http://doi.org/10.1145/2047196.2047246>
38. P. Viviani and C. Tenuolo. 1982. Trajectory Determines Movement Dynamics. *Neuroscience* 7, 2: 431–7.
39. Gerold Wesche and Hans-Peter Seidel. 2001. FreeDrawer: a free-form sketching system on the responsive workbench. ACM Press, 167. <http://doi.org/10.1145/505008.505041>
40. Eva Wiese, Johann Habakuk Israel, Achim Meyer, and Sara Bongartz. 2010. Investigating the learnability of immersive free-hand sketching. *Proceedings of the Seventh Sketch-Based Interfaces and Modeling Symposium*, Eurographics Association, 135–142.
41. Baoxuan Xu, William Chang, Alla Sheffer, Adrien Bousseau, James McCrae, and Karan Singh. 2014. True2Form: 3D Curve Networks from 2D Sketches via Selective Regularization. *ACM Trans. Graph.* 33, 4: 131:1–131:13. <http://doi.org/10.1145/2601097.2601128>
42. Youyi Zheng, Han Liu, Julie Dorsey, and Niloy J. Mitra. 2016. SmartCanvas: Context-inferred Interpretation of Sketches for Preparatory Design Studies. *Computer Graphics Forum*, Wiley Online Library, 37–48.
43. Tilt Brush by Google. Retrieved January 8, 2017 from <https://www.tiltbrush.com/>
44. PaintLab VR - Paint and Sculpt in Virtual Reality. *PaintLab VR*. Retrieved September 20, 2016 from <http://paintlabvr.com>
45. DesignSpace. Retrieved September 20, 2016 from <http://www.designspacevr.com/>
46. Tvorì—Anyone Can Tell a Story. Retrieved September 20, 2016 from <http://tvori.co>
47. Virtual Art Sessions. Retrieved September 20, 2016 from <https://virtualart.chromeexperiments.com/>

<sup>i</sup> © Google—<https://developers.google.com/web/showcase/2016/art-sessions>. Used under Creative Commons 3.0 Attribution.

<sup>ii</sup> © Gravity Sketch—<http://www.gravitysketch.com>. Used with permission.