

WieldingCanvas: Interactive Sketch Canvases for Freehand Drawing in VR

Xiaohui Tan^{*†}
Capital Normal University, Beijing
Beijing, China

Zhenxuan He^{*}
Institute of software, Chinese
Academy of Sciences
Beijing, China

Can Liu
City University of Hong Kong
Hong Kong, China

Mingming Fan[‡]
The Hong Kong University of Science
and Technology (Guangzhou)
Guangzhou, China

Tianren Luo[§]
Institute of Software, Chinese
Academy of Sciences
Beijing, China

Zitao Liu
Guangdong Institute of Smart
Education, Jinan University, China
Guangzhou, China

Mi Tian
TAL Education Group
Beijing, China

Teng Han^{¶||}
Institute of Software, Chinese
Academy of Sciences
Beijing, China

Feng Tian[¶]
Institute of Software, Chinese
Academy of Sciences
Beijing, China

ABSTRACT

Sketching in Virtual Reality (VR) is challenging mainly due to the absence of physical surface support and virtual depth perception cues, which induce high cognitive and sensorimotor load. This paper presents WieldingCanvas, an interactive VR sketching platform that integrates canvas manipulations to draw lines and curves in 3D. Informed by real-life examples of two-handed creative activities, WieldingCanvas interprets users' spatial gestures to move, swing, rotate, transform, or fold a virtual canvas, whereby users simply draw primitive strokes on the canvas, which are turned into finer and more sophisticated shapes via the manipulation of the canvas. We evaluated the capability and user experience of WieldingCanvas with three studies where participants were asked to sketch target shapes. A set of freehand sketches of high aesthetic qualities were created, and the results demonstrated that WieldingCanvas can assist users with creating 3D sketches.

CCS CONCEPTS

• **Human-centered computing** → **Virtual reality**.

^{*}Both authors contributed equally to this research.

[†]Also with Institute of Artificial Intelligence Education Capital Normal University.

[‡]Also with The Hong Kong University of Science and Technology.

[§]Also with School of Computer Science and Technology, University of Chinese Academy of Sciences.

[¶]Also with School of Artificial Intelligence, University of Chinese Academy of Sciences.

^{||}Corresponding author. Email: hanteng@iscas.ac.cn

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 the author(s) 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](https://permissions.acm.org).

CHI '24, May 11–16, 2024, Honolulu, HI, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0330-0/24/05...\$15.00

<https://doi.org/10.1145/3613904.3642047>

KEYWORDS

freehand drawing, interactive canvas, two-handed interaction, VR

ACM Reference Format:

Xiaohui Tan, Zhenxuan He, Can Liu, Mingming Fan, Tianren Luo, Zitao Liu, Mi Tian, Teng Han, and Feng Tian. 2024. WieldingCanvas: Interactive Sketch Canvases for Freehand Drawing in VR. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3613904.3642047>

1 INTRODUCTION

The recent advancement of Virtual Reality (VR) technologies has facilitated creativity expression in a wide range of artistic applications including painting, drawing, 3D modeling, sculpting, and animation. The immersive visuals and spatial input capabilities have reshaped the ways how designers express and communicate ideas, but meanwhile also induced new challenges as a result of extensions to the spatial interaction. For instance, freehand sketching is among designers' favorites for product ideation and allows them to create various types of strokes to form accurate and aesthetic shape projections [39]. However, it is challenging to sketch freehand in 3D space which often results in the lack of precision and control, due to the absence of physical surface support and virtual depth perception cues, as well as high sensorimotor demand [5, 10, 25, 36].

Early work supported freehand sketching in 3D with auxiliary hardware for sensing user input or providing haptic feedback [8, 15, 20, 21]. Recent research efforts investigated ways to beautify, neaten, and auto-correct strokes by providing visual guidance to allow users to sketch more precise lines and curves in 3D [5, 25, 27, 33, 38, 39]. As opposed to previous work on curve beautification and correction that attempt to automatically predict the desired shapes and apply corrections based on priors, we propose to let users indicate their intent more explicitly by sketching freehand with facilitated precision and control. The interactions are inspired by craft activities, in which users create precise shapes




Figure 1: Concept design of WieldingCanvas: Invoke a spinning canvas with one hand and touch the canvas with the other hand to draw a circle.

such as replicable, symmetrical, Spin Art, cutting or sculpting on papers, whiteboard, or clay that are manually rotated or spinning automatically on a potter’s wheel (Figure 3). In these activities, precise and controlled actions are facilitated via creating regular movement of the painting board to help maintaining stability of the drawing hand, that often engage cooperations of both hands. Spin Art, for instance, involves rotating a canvas or paper while applying drops of paint to create various patterns and designs. Similarly, the use of rotational motion can aid in creating precise and symmetrical shapes [2]. Although the specific techniques may vary, the key principle is using regular movements to achieve precision and replicate controlled forms.

This paper presents WieldingCanvas, an interactive VR sketching canvas with the aim to enable users to sketch freehand in 3D with improved precision while maintaining explicit user control over the result instead of doing automatic correction. The system provides users with manipulable virtual canvases and two novel strategies, including i) the users can indicate the type of geometric shape they desire through non-drawing hand gestures that spawn planar canvases with motions that generate regular shapes (Figure 1), and ii) the users can draw planar curves on a virtual planar canvas, then deform them into 3D spatial curves with two-hand gestures. In this paper, we first describe the design of WieldingCanvas, capturing innovative ways to support motions and transformations of the canvas that facilitate users’ sketching actions. A set of mid-air gestures are defined to invoke WieldingCanvas with desired motions (e.g., rotating, oscillating, swinging), or conduct target transformations (e.g., curving, twisting) on sketched canvases. The design is examined and refined by a group of five interaction designers in VR and AR domains. We then present a prototype application of WieldingCanvas, that demonstrates the coordination between canvas manipulations and sketch actions to empower users to create freeform yet refined lines and curves in VR. Interface design factors including the appearance, visual and audio cues, haptic responses, affordance, and management of the canvases are carefully considered and defined.

We present three empirical studies to evaluate whether WieldingCanvas supports users on the following 3 criterion: drawing precision (error compared to a target, smoothness), completion time and task load. Additionally, a usability study is provided. Participants with few VR sketching experience were recruited and asked to sketch primitive lines and curves, target 3D shapes, and free

creations, respectively. The results indicated that WieldingCanvas can support expressive creation of lines and curves in 3D. The novel functionalities improve drawing precision and smoothness, and lead to lower frustration in performing the task compared to the baselines, i.e., freehand 3D sketching and sketching on non-manipulable virtual canvas. Although it costs more time in general to complete the drawings using WieldingCanvas, most participants still preferred using the interactive canvases to assist them to draw.

The paper made the following contributions: i) WieldingCanvas as a novel assistive 3D sketching technique that cooperates canvas manipulations and freehand drawing activities; ii) the design and implementation of interactive canvases that provide motions and transformations to facilitate users’ sketching activities with improved precision and smoothness; iii) the evaluation of WieldingCanvas that indicates how sketch operations combined with canvas manipulations can advance users’ freehand 3D sketch capabilities in VR.

2 RELATED WORK

WieldingCanvas is designed to empower users’ freehand sketching capabilities in VR via integrating canvas manipulations and sketch actions. We draw inspirations from previous work on exploring sketching interfaces in VR, creative interactions with both hands, and the use of physical and virtual surfaces in VR/AR.

2.1 VR Sketching Interfaces

As VR technologies are getting more prominent, commercial sketching applications such as TiltBrush [14] and GravitySketch [33] are widely available and have captured users’ imagination on novel ways to sketch in 3D. Though very popular, it is known that VR sketching is non-trivial as people are used to sketching in 2D plane [37], and when it comes to the 3D space, people face challenges such as depth perception errors and the absence of a physical surface [5, 9]. These limitations result in less accurate drawings but higher cognitive and sensorimotor load [25]. Addressing the challenges in VR sketching, studies such as those by Machuca et al. [24] delve into the selection of appropriate interaction devices and techniques, considering user needs and the unique affordances of different tools in 3D sketching environments. Wiese et al. [36] investigated the learnability of 3D sketching and found that with training, sketch quality was improved but time cost remained unchanged, and the shape of the target seemed to have no influence on sketching performance. The work raised an interesting question to what extent do the higher error-proneness and higher sensorimotor demands cause the loss of performance quality for 3D objects respectively. Arora et al. [5] further looked into how physical constrains, visual guidance, orientation, and scale affected users’ 3D drawing capabilities and pointed out that including physical surfaces and visual guidance is efficient in reducing the loss of drawing precision.

Recently it has been of special interest to researchers to look into how to empower users’ sketching abilities in VR. For instance, Machuca et al. [25, 27] leveraged the design of snapping planes, beautification trigger points, and a set of visual guides to help users improve shape-likeness and stroke precision. Yu et al. [38] explored a 3D optimization method that helps users create automatically connected and predicatively smoothed curves. Singh et

al.[32] explore the integration of traditional French curves with contemporary computer-aided design. Grossman et al.[16] presents a novel interface for directly manipulating curves in both 2D and 3D environments. Fung et al.[13] present Kinematic Templates as innovative end-user tools for content-relative cursor manipulations, designed to assist in the creation and manipulation of 2D visual compositions. Yu et al. [39] provided product designers with visual references (i.e, scaffold) and auto-correction to draw accurate and aesthetic shape strokes. In addition to these auto-correction, beautifying, and neatening approaches, researchers investigated ways to exploit physical surfaces during drawing. Wacker et al. [35] studied the feasibility of using physical objects as tangible references to direct 3D sketching. Drey et al. [11] and Henrikson et al. [17] used additional devices such as tablet to assist drawing, where users stick to the traditional and familiar tablet + pen input method while viewing the sketches in VR. While much of the existing research focuses on VR environments, several studies in non-VR settings offer valuable insights. For example, the work by Ikeda et al. [19] employs canvas rotation to assist users in quickly sketching fluid targets. This approach highlights the potential benefits of canvas manipulation techniques in enhancing user efficiency. The importance of such advancements is underlined by Machuca et al.[26]’s work, which presents a comprehensive framework for evaluating these emerging tools and methodologies in the context of 3D immersive sketching.

Inspired by creative activities in real life, we present an alternative 3D sketch input metaphor for users to integrates canvas manipulations into drawing actions, in order to facilitate the creation of 3D lines and curves in VR.

2.2 Two-Handed Creative Interactions

Two-handed activities are natural for human beings and can be categorized based on how the brain coordinates two-sided or bilateral movements, including symmetrical, alternating, and using a more active movement side together with a stabilizing side [28]. In daily lives, we use two-handed movements to create artifacts such as paper cutting, making pottery, and drawing with pen and paper, and previous research showed that it is beneficial to borrow the "feel" from a physical scene, and use two-handed actions in innovative ways to exploit the dexterity from both hands in interactive sketching systems. Bae et al. [6, 7] designed a 3D curve sketching system that integrated the affordance of pen and paper with several novel features. In their system, a user’s non-preferred hand was used to press buttons to trigger clutching functions. Kim et al. [22, 23] creatively used a non-drawing hand to provide hand posture information to sketch 3D handheld products and to generate approximate 3D shapes as well as sketching scaffolds.

Some other notable works were based on touchscreen tablets, where touch gestures were combined with pen-based sketch input to provide additional affordance. For instance, Henrikson et al. [18] provided rich interactions for depth specification in sketching stereoscopic storyboards by using a non-drawing hand to directly manipulate stereoscopic visualizations while creating content. Pfeuffer et al. [29] explored two-handed interactions on tablets that unified touch gestures from one hand with 2D pen commands from the other hand for 3D sketching and manipulation. Our work was

inspired by these works and explored the design space of cooperating both hands for expressive line and curve sketching in VR with the assistance of interactive canvases. Different from previous work, the hand actions in WieldingCanvas involve 1- the non-drawing hand invokes a virtual functional canvas (e.g., rotating, oscillating, swinging) whereby the drawing hand sketches on it, and 2- using both hands to transform a canvas (e.g., bend, twist, stretch) with sketched lines or curves. Grossman et al. [16] present an interface for far more direct manipulation of curves using a specialized two-handed controlled input device. Verhoeven et al. [34] introduced a prototype that allows users to deform a mesh using their fingers on both hands.

2.3 The Use of Surfaces in VR/AR

Compared to sketching in 2D, drawing in 3D (e.g., VR) introduces a higher cognitive and sensorimotor workload. To alleviate the drawing burden in 3D, researchers have explored various forms of surfaces, including physical and virtual ones, and integrated them as part of VR/AR sketching interfaces. One straightforward idea is to use tablets. Xin et al. [37] used a tablet and pen to create 3D sketches on top of a physical napkin. Arora et al. [4] combined 2D and 3D sketching with the help of a tablet and for AR applications. Drey et al. [11] took a similar approach and investigated the combination of the 2D sketch on a tablet and the 3D sketch in mid-air for VR. On the other hand, creating and manipulating virtual surfaces is useful as it provides visual and sketching references. For instance, Bae et al. [7] supported automatic sketch plane selection, extruded surface sketching, and freeform surface sketching in their system. Schholne et al. [30] supported users to mark 3D space and create organic 3D shapes directly from hand motions. Similarly, Fuge et al. [12] and Kim et al. [22] asked users to use their non-drawing hands to create surfaces in the virtual space.

Our work also requires virtual surfaces and uses them as interactive canvases to support primitive drawings and shape creation. An earlier work on 3D stroke painting in a 2D viewport window by Schmid et al. [31] shared a similar idea about using shaped canvas. They proposed a canvas defined by implicit surfaces of approximate 3D proxy, which allowed users to directly paint 3D strokes using the canvas. Compared to their work, WieldingCanvas not only supports shaped canvas but also allows for integration of canvas manipulation and sketch actions in VR.

3 DESIGNING INTERACTIVE SKETCH CANVASES




Figure 2: When sketching, the drawn lines are accurate in the user’s field of view, but when viewed from another angle, it will find that the deviation is very large.

Freehand sketching in VR can be difficult for drawing regular shapes, such as straight lines or circles, due to users' lack of experience or skill in controlling depth with arms and hands in the air [3]. Figure 2 illustrates the problems for freehand sketching in VR from real examples, that the lines and curves are not smooth and cannot be aligned to single planes. We aim to improve the sketching precision while maintaining full user control in the hand drawing process, as opposed to having a system "beautifying" the sketch automatically.




Figure 3: The creative inspiration of the WieldingCanvas: (a) drawing a circle by moving the paper instead of the pen. keep pen still and rotate. (b) paper-cutting by folding and turning the paper instead of moving the scissors. keep right scissors still and rotate. (c) Spin art on white board.¹ (d) clay on potter's wheel.²

The design of WieldingCanvas is inspired by the underlying principle common in various handcrafting processes (Figure 3), where controlled manipulation of the medium results in enhanced outcomes. This principle is exemplified in pottery, where regulating clay deformation on a rotating base is key. In cake decoration, chefs rotate the plate for precise cream curvatures; and in paper cutting, where controlled rotation and folding of the paper aid in achieving symmetry and precision. Inspired by this principle, our solution for addressing 3D hand-drawing inprecision is centered around introducing virtual canvases that can be repeat regular motions or deformed for users to sketch on. Based on this framework, we expanded standard motion into rotational, linear, and pendulum motions, and elaborated deformation to include stretching, bending, twisting, and folding as central components of our design. By making use of the simulated physical properties of the canvases, we introduce a new approach to facilitate the manual creation of regular shapes or smooth 3D curves. We envision WieldingCanvas as having the capability to integrate seamlessly into existing drawing processes, thereby boosting performance in sketch-based activities.

3.1 Hand Action Sequences

When incorporating canvas manipulations with users' drawing actions, it is crucial to figure out suited hands' operations, e.g., in a simultaneous, sequential, symmetrical or alternating manner. First, canvas manipulations could happen after the drawings to deform sketched lines or curves, whereby the actions could be straightforward and get both hands involved. Second and what is more challenging, is that manipulating the canvas for drawing with both hands acting simultaneously, like sketching or cutting while rotating the paper (Figure 2), would require the non-drawing hand and

the drawing hand move synergistically while maintaining accurate and steady. However, unlike real-life scenarios, simultaneous spatial operations of both hands can be problematic. During the design exploration process, we discovered that it was challenging to simultaneously operate the canvas with one hand while maintaining the control over the brush with the other hand. Specifically, due to the lack of physical constraints, the action of one hand (e.g., rotating the canvas with a controller) affects the stability of the other (i.e., the drawing hand) in the mid-air, which might be unnoticeable to users themselves.

To verify if this problem exists, we carried out an informal pilot study and invited 6 participants (2 female, avg. age = 24.3) from a local university who had no experiences in VR before. We implemented an application in VR that replicated the scene of drawing circle via rotating the paper. We asked the participants to rotate the virtual canvas with the controller in their non-drawing hand, and meanwhile, touch the canvas with the other controller (as the brush) and hold steady using their drawing hand while standing. We recorded the results and found that the passive shaking of the drawing hand due to the motions of the non-drawing hand was inevitable and affected the outcome, resulting in swirling circles, as illustrated in Figure 4.




Figure 4: Due to the lack of physical constraints, rotating the canvas affects the stability of the drawing hand in the mid-air.

To avoid such a problem, we suggest the two hands act sequentially instead of simultaneously, meaning that users could rotate the canvas first and then draw the circle on the spinning canvas, as shown in Figure 7. In this way, users can maintain the stability of the drawing hand by keeping body posture still. There could be feasibility where users manipulate the canvas and sketch on it, both with the drawing-hand. However, to make it simple and the UI function less confusing, we decided to go for the way using both hands, and let the non-drawing hand to invoke the canvas, and the drawing hand to control a virtual brush and sketch.

Based on these hands' operations, WieldingCanvas is designed to include sequential hands' actions to coordinate canvas manipulations and drawing actions, and simultaneous hands' actions to transform the canvas with sketched lines and curves, both to enhance users' freehand sketching capabilities. The design of WieldingCanvas is elaborated on in the following.

3.2 Canvas Manipulations and Drawings

After deciding on the hand manipulation sequence, the next step is to design and prototype the canvas manipulation for supporting

¹<https://images.squarespace-cdn.com/content/v1/5e0e1aed65066232ebd5e14d/1647213685965-ONQE2W8VR60S2A10ZRBS/SpinningArtRobotInAction2.JPG>

²https://keyassets.timeincuk.net/inspirewp/live/wp-content/uploads/sites/8/2022/04/CLL350.ints_ceramic.0037-2-920x764.jpg

drawing. We explore the use of two types of basic physical manipulations for paper: *motion* and *deformation*. Although the exhaustive lists of both types of operations can be infinite, we chose a few representative operations for testing our approach with a proof-of-concept prototype. Our design of WieldingCanvas supports the activation of three types of canvas movement before drawing on it: *Transport*, *Rotate* and *Oscillate*. Users could also deform the canvas after drawing on it to add 3D curvature. Four types of deformations are supported: *Fold*, *Bend*, *Stretch* and *Twist*. Below we describe in detail how the canvas motions are created and activated via non-dominant hand gestures before drawing, as well as how canvas deformation can be performed with both hands, akin to in the physical world, and the curvatures it adds to the drawings.

3.2.1 Invoking Gestures. As said, users initially used their non-drawing hand/controller to manipulate the canvas. However, when starting with a blank canvas in VR, users had to first invoke the canvas through a button press and menu selection, followed by positioning and orienting the canvas for drawing. This process was found to be inefficient and required too many redundant steps. To address this issue, an optimized gestural approach was implemented for canvas invocation. Users can now perform a simple gesture, such as circling or waving their non-drawing hand (with the controller’s thumb button pressed) in mid-air to initialize the canvas with desired functions, such as spinning or oscillating. The system then recognizes the gesture and the corresponding canvas appears in the scene (see Figure 5). This new approach reduces the number of steps required and improves the efficiency of canvas invocation.



Figure 5: Four canvas invocation gestures of the motion canvas, and use the displacement of the handle directly as the input of the gesture action: a) circle action. b) straight action. c) curved action. d) fold action.

The gesture recognition is based on the movement trajectory of the hand. When a user presses the invoking button, visual particles are displayed along with the controller, and the system creates a trailing effect that signifies the recording of the hand motion trajectory for gesture recognition. To streamline the process, enhance real-time computational efficiency, and acknowledging that not all points need spatial information, we opted to project point data onto a plane. We provide two plane options for users: one that always faces the user and another based on the fitted plane of all points. The user-facing plane, preferred by novices, offers visual effect that engages them, making it easier for beginners. The fitted plane option, with its three-dimensional coordinates and orientation, allows for control over the canvas’s final position. This approach balances simplicity with user preference and functionality. The recognition process relies on several key feature values, which

encompass the radius of the fitted circle, linear correlation, the presence of a V-shaped pattern, detection of a W-shaped bottom, and the sum of angles. The first four feature values are obtained through well-established standard algorithms. In contrast, the ‘sum of angles’ is a unique method employed to estimate the cumulative radians of the fitted shape. The formula for calculating sum of angle is as follows:

$$sum_{an} = \sum_{i=1}^{n-1} Angle(P_c \vec{P}_i, P_c \vec{P}_{i+1})$$

where P_i denotes the sampling point, while P_c refers to the center of the fitted circle. The adjacent vectors that connect the hand trajectory sampling point and the trajectory circle are represented by $P_c \vec{P}_i$ and $P_c \vec{P}_{i+1}$ respectively (Figure 6a).

Once the trajectory is recognized as an invoking gesture, the corresponding canvas gets created and its initial parameters, including its position, orientation, and speed, are determined by the gestural trajectory. More detailed descriptions are included in the following subsections.




Figure 6: Trajectory of controller/hand: a) Circle. b) Straight line. c) Arc. The red dot represents the starting point (P_0), and the green dot represents the ending point (P_1).

3.2.2 Circle. With WieldingCanvas, circles are drawn by first invoking a rotating canvas via users tracing their non-drawing hand in an anticlockwise circular pattern, and then painting at a single point onto the spinning canvas with the drawing hand (Figure 7). This gives users novel 3D sketching experiences with the ability to rapidly draw a circle with a higher level of control and precision. WieldingCanvas initializes the spinning canvas and sets its position based on the invoking gesture. Specifically, the center of the canvas is determined by the center of the fitted circle. The initialization speed and rotational direction of the canvas are calculated based on sum_{an} , as listed in the previous subsection. Moreover, the canvas rotating speed increases with the number of turns made by the user. While a faster rotating speed allows for faster drawing, it may be more challenging to maintain the precision. This design enables users to set their preferred rotating speed, providing skilled users with the ability to improve their efficiency.

The canvas keeps rotating after the controller leaves and decelerates with a coefficient of friction of 0.01. Users can also rotate an existing canvas in the scene by touching the controller to the canvas and perform the circling actions.

3.2.3 Straight Lines. We followed a similar approach and designed the interaction to support users to draw straight lines on the canvas. Users perform a horizontal waving gesture from left to right, or




Figure 7: The 4 steps of drawing circle with spinning canvas: a) Perform a circle action by left hand. b) Generating canvas in motion. c) Drawing by right hand. d) Completed circle.

vice versa with the non-drawing controller and the trigger button pressed. To calculate the real-time position of the canvas after users' action, we use the following equation:

$$P = P_0 - P_0 \vec{P}_1 \cdot e^{-bw} \cos(t\sqrt{w^2 - b^2w^2})$$

This equation is used when the system detects that the trajectory is sufficiently linear (i.e., larger than 0.5) and $|sum_{an}| < 50^\circ$, and a virtual canvas fades in and starts to oscillate. The start and end points of the trajectory define the two ends of the oscillating motion, while P_0 and P_1 are the starting and ending points of the gesture (Figure 6b), respectively. The spring slope and friction coefficients, w and b , respectively, are determined empirically and affect the speed of simple movement.

Users can then put and hold the drawing controller still on the canvas via the other hand, to draw a straight line, as shown in Figure 8. To work with an existing canvas, users first drag the canvas either horizontally or vertically with one hand and release it to make the canvas move back and forth along the dragging direction, like a spring. Besides straight lines, such a method allows users to purposely move the drawing hand, e.g., slightly in the direction perpendicular to the canvas's moving direction at each inflection point, to draw parallel lines or reciprocating lines to represent a surface.




Figure 8: The 4 steps of drawing straight line with moving canvas: a) Perform a straight action by left hand. b) Generating canvas in motion. c) Drawing by right hand. d) Completed straight line.

3.2.4 Curves. A similar approach was applied to assist users with drawing curves. First, users swing the non-drawing hand to invoke the canvas and make it do pendulum movement by gravity. The canvas appears and starts moving at the end of the gesture trajectory, and performs a pendulum motion on the plane where the fitting arc of the trajectory is located. The real-time angular acceleration of motion can be calculated as:

$$\alpha = \frac{Distance(O, P_1) \cdot G}{Distance((O.x, P_1.y, O.z), O)^2}$$

where G is the gravitational acceleration, and the movement acceleration α is determined by G . The movement velocity has an initial value of 0 and changes every frame according to the α . O is the center of the fitting circle of the arc, and P_1 is the endpoint of the gesture trajectory (Figure 6c). Then users keep their drawing hand static on the canvas. In this way, a smooth curve can be drawn, as shown in Figure 9. Similarly, to work with an existing canvas, users can press the button and drag the corner of the canvas towards a position. Releasing the button will let the canvas take the initial position as the drawing point, and the current position as the starting point for a pendulum movement.




Figure 9: The 4 steps of drawing smooth curve with moving canvas: a) Perform an arc action by left hand. b) Generating canvas in motion. c) Drawing by right hand. d) Completed curve.

3.2.5 Repetitive Shapes. The design of the foldable canvas is inspired by the traditional Chinese folk art of paper-cutting. Folding operation helps create symmetrical shapes, which often appear in drawings (Figure 10). First, users create a canvas folded in two or four equal parts and then draw a pattern on the folded canvas. Unfolding the canvas forms two or four copies of the same pattern, which results in a fine drawing with symmetric patterns. Similar to the previous ones, users perform a gesture to indicate the canvas folding command, where a 'V' or 'W' shape gesture represents a canvas being folded once or twice. To work well with the folded canvas, it prefers users to have good imagination of the final effect based on the pattern to sketch. The folded canvas simplifies user operations and extends the ability to draw symmetric shapes.

3.3 Canvas Transformations

Transformable canvases enable users to begin with basic shapes and then apply deformation operations on them, such as bending, stretching, and twisting operations on canvases, to generate complex 3D curves. Three types of transformable canvases are designed to help users draw complex curves or planes with corresponding operations.

3.3.1 Bend Operation. Users can bend the canvas to create 3D curves. The design is derived from actions to bend a real paper. As shown in Figure 11b, the user grasps the two edges of the canvas and rotate the hands in opposite directions with sketched 2D lines. The rotation on Roll axis of the handle controls the degree of bending.




Figure 10: There are two ways to call a foldable canvas, with drawing a 'V' or 'W' to invoke 1-fold or 2-fold canvases, respectively. 1-fold steps: a) Perform a 'V' action by left hand. b) Drawing by right hand. c) Unfold the canvas. 2-fold steps: d) Perform a 'W' action by left hand. e) Drawing by right hand. f) Unfold the canvas.




Figure 11: The three deformation methods require both hands to operate at the same time: a) Original. b) Bend Operation. c) Stretch Operation. d) Twist Operation.

The curved shape is adjusted using a cubic Bézier curve formula, as shown in the following:

$$P = P_l \cdot (1-t)^3 + (P_l + \vec{left}) \cdot 3(1-t)^2 t + (P_r + \vec{right}) \cdot 3(1-t)t^2 + P_r \cdot t^3$$

where \vec{left} and \vec{right} represent the direction vector after rotation on the Roll axis. P_l represents the grasping position of the left hand and is also the starting point of the Bézier curve. P_r represents the grasping position of the right hand, which is also the end point of the Bézier curve. t is the parameter that determines the position of a point P on the curve, ranged from 0 to 1. When t is 0, P is at the start of the curve (at P_l), and when t is 1, P is at the end of the curve (at P_r). The value of t between 0 and 1 determines the position of P on the curve between P_l and P_r (Figure 12a).

Users may want to draw curved strokes like a 'S' shape in 3D space. In these cases, the bending canvas may produce a curve with desirable perfect curvature because it is easier to operate with the canvas than sketching on the canvas. Meanwhile, with this proposed method, not only symmetric shapes but also asymmetric shapes can be generated. Users only need to rotate their hands to different degrees to create asymmetric curved shapes following the rule of the Bézier curves.

3.3.2 Stretch Operation. In the stretching mode, the canvas can be stretched so that elastic curves can be created with less effort. Canvas can be extended or shortened in one direction (Figure 11c). With stretching operation, the user can hold the two edges of the canvas, and then according to the distance variation between the two controllers, the curves on the canvas can be stretched or compressed accordingly. When the two edges of the canvas are held, the global deformation is applied. When the middle parts are held,




Figure 12: Canvas Transformation and Deformation: a) Bend Deformation. b) Stretch Deformation. c) Twist Deformation. The point P_0 represents any point on the canvas or stroke before deformation, and P represents the same point after deformation. Canvas transformations are used to adjust the shape and scaling of the canvas or stroke, allowing for different forms of deformation to be applied.

only the local part between the controllers is deformed without affecting other parts. Such a global and local fine-tuning mechanism enables users to modify the shape of curves with high flexibility. The calculation of the shape changing happens in the local space as following:

$$P = P_0 + \vec{OP}'_0 \cdot (\text{rate} - 1)$$

where rate is obtained based on the distance between the two controllers, which determines the degree of boost or compression. OP'_0 represents the origin of the stretching direction, P_0 represents the initial point before deformation, and P'_0 is projection of this point in the stretching direction (Figure 12b).

3.3.3 Twist Operation. Twist operation allows the canvas to rotate in a spiral pattern and thus extend lines or curves to the third dimension (Figure 11d). Users can create spiral 3D strokes that are difficult to create with existing VR drawing tools. With this operation, users just need to hold both edges of the canvas and twist them against each other to generate a spiral 3D curve or a surface. This operation allows users to generate centrosymmetric three-dimensional shapes by twisting the canvas spirally according to the angle of the controller. The deformation is also calculated in local space, and the point on mesh is represented as following:

$$P = (x_0 \cos(z_0 (An_l - An_r)) - y_0 \sin(z_0 (An_l - An_r)), \\ x_0 \sin(z_0 (An_l - An_r)) - y_0 \cos(z_0 (An_l - An_r)), z_0)$$

The angles An_l and An_r are obtained from the variation of the Pitch axis angle of the controllers, where each hand controls one input. The degree of twist increases as the difference between An_l and An_r increases. The variable z_0 is mapped from 0 to 1 in the local space, and the point (x, y, z) is deformed such that $z = z_0$ remains unchanged before and after deformation. The deformation of the x and y components of the point can be seen as a two-dimensional vector rotation formula, but the rotation angle is determined by $z_0 (An_l - An_r)$. The linear change of z_0 causes the mesh to change spirally in three-dimensional space after deformation (Figure 12c).

3.4 WieldingCanvas Interface Design

WieldingCanvas leverages the advantage of enabling sketching with canvas manipulations, with which users can create various accurate

and smooth lines and curves in 3D space using the VR controllers. Besides the functions of the canvases, their use in VR requires significant extra design considerations to be effective, efficient and easy to learn.

3.4.1 Canvas Management And Accessibility Features. The canvases provide the functions to sketch lines, circles, and curves, and in the actual sketching activities, users may often need to assemble the sketches and adjust their positional relationship. This brings up the requirement for snapping the canvases to facilitate user manipulation of the sketches. For instance, when drawing a group of parallel wheels, parallel snapping to align the canvases can reduce users' workload.

In actual sketching activities, users do not need the canvas to visually appear all the time. We designed a method for users to toggle the showing and hiding of the canvas. When users finishes working on a canvas, they can hide the canvas by pressing Y when it is selected. In the hiding state, only the strokes on the canvas are visually available. Contacting the hidden canvas shows a red dot in the center so that users are aware of the position of the canvas.

WieldingCanvas offers button-less interaction for basic drawing, enabling immediate sketching upon touching the canvas to reduce response time. The advanced functions, compatible with Tilt Brush UI, include button-based drawing beyond the canvas, rich brush customization, and common editing features like copy, undo, and redo.

3.4.2 Visual, Auditory, and Vibration Feedback. The goal of the vibration feedback is to meet challenges of 3D freehand sketching which lacks adequate feedback. When users touch the canvas with the non-drawing controller, a slight vibration is provided to remind them that the canvas is ready to be used (Figure 13). There is no vibration feedback when the drawing controller touches the canvas. Only when users press the trigger button to indicate they are ready to sketch, the controller gives strong vibration to prompt the beginning of drawing. When the drawing controller is used to sketch on the canvas, users feel a weak friction shock, which is designed to provide a similar experience of drawing on a piece of paper.

Besides, visual feedback is designed to indicate different sketching stage. When the non-drawing controller touches the canvas, the canvas gets yellow highlighted around its edge as a prompt. The color of the brush turns blue when the drawing controller touches the canvas, and it turns pink when it get away from the canvas. Auditory feedback is provided when users select the specific function of the canvas. For instance, when users sketch on the canvas with the drawing controller, users hear a faint friction sound to increase users' immersion experience.

4 EVALUATION

Our evaluation of WieldingCanvas meticulously examines two key aspects [26]: 1) the assessment of sketching artifacts, and 2) the evaluation of sketching tools. These evaluations were systematically carried out in Study 1 and Study 2, respectively.

Study 1 explores whether enabling canvas manipulation in VR can improve the precision of freehand drawing for basic shapes. The hypothesis for Study 1 is that the ability to manipulate the




Figure 13: The feedback of the canvas is divided into non-drawing controller feedback and drawing controller feedback: a) The state of the canvas when it is not in use. b) When the left hand touches the canvas, the edge of the canvas will be highlighted and vibrated at the same time. c) The brush of the drawing controller touches When painting the canvas, it vibrates a bit, while during painting it provides intermittent vibrations and constant sounds to simulate friction.

canvas will lead to increased precision and quality in drawing basic shapes.

Study 2 examines whether WieldingCanvas can enhance the quality of hand-drawn complex 3D objects. It also seeks to understand the tool's versatility in supporting open-ended hand-drawing tasks. The hypothesis for Study 2 is that WieldingCanvas, by facilitating more natural and intuitive drawing movements, will not only improve the quality of complex 3D hand-drawn objects but also provide a effective tool for different drawing tasks.

Study 1 and Study 2 are interconnected, with each building upon the other in a progressive assessment of our technique's capabilities. Study 1 centers on the precision of replicating basic shapes, testing foundational elements like ease of use and precision. Study 2, building upon Study 1's outcomes, tackles more complex 3D drawing tasks to evaluate the technique's versatility and effectiveness in diverse, intricate scenarios. Together, these studies form a comprehensive evaluation, demonstrating the technique's efficacy from basic to advanced sketch tasks.

4.1 Apparatus and System

The hardware used is Meta Quest 2 and our system is ingeniously developed from the open-source version of Tilt Brush [1], utilizing C# and Unity along with the Oculus and SteamVR SDKs. The open-source version of Tilt Brush, while excluding certain premium resources, maintains a level of functionality that closely parallels that of the commercial version. As a direct fork of Tilt Brush, our platform inherits the entire spectrum of its capabilities. We've conducted extensive testing to ensure its robust performance on the Quest2 platform, which includes the process of APK packaging and installation. To further refine our understanding of user interaction with the system, we employed the Virtual Desktop streaming tool in Study 1&2. To facilitate this, our system operated on a high-performance laptop with an AMD Ryzen9 5900HX CPU and NVIDIA GeForce RTX 3070 Laptop GPU, while all display and interactive elements were seamlessly executed on the Quest headset. This setup provided an efficient and comprehensive means to observe user engagement with the system.

4.2 Study 1: Evaluating Sketching Artifacts

We conducted an empirical user study to assess the capacity and user experience of WieldingCanvas, focusing on it supports users to leverage canvas manipulations to create lines, curves, and shapes in VR. We conducted comparative analyses between WieldingCanvas, OnlyCanvas, and TiltBrush [14]. OnlyCanvas, a derivative of WieldingCanvas, exclusively employs a brush for drawing on the virtual canvas, eschewing the use of other techniques. The inclusion of OnlyCanvas in our study was intended to isolate and observe the effects attributable solely to the virtual canvas. TiltBrush represents a commercial VR system distinguished by its capability to facilitate painting directly in mid-air.

4.2.1 Design. The study was conducted using a within-subject design, where each participant was tested across three different technologies (*WieldingCanvas*, *OnlyCanvas*, *TiltBrush*) and five different tasks (*circle*, *line*, *bow*, *tilde*, *ellipse*). These geometries are designed to be drawn in a single, continuous motion. Furthermore, all selected geometries can be replicated using each of the tested techniques, ensuring a consistent basis for comparison. The key distinction between drawing *circles* and *ellipses* lies in the need for sequential utilization of motion and deformation functions when creating *ellipses* with WieldingCanvas. The situation is similar to *tilde*. The sequence in which technologies were presented was counterbalanced across participants using a Latin Square design. And each task was performed three times by each participant. Tasks appeared in random order, and participants were required to complete all tasks associated with one technology before moving on to the next. Each of the 5 tasks was repeated three times for each technology, resulting in a total of 45 strokes per participant.

4.2.2 Task. A group of five basic geometries were selected that included *circle* (radius=21.7cm), *line* (length=52.8cm), *bow* (with=54.1cm, height=9.5cm), *tilde* (with=52.2cm, height=23.7cm), *ellipse* (with=54cm, height=36cm), as shown in Figure 14. During the drawing tasks, the task targets remained visible, allowing participants to continuously track the target while sketching.



Figure 14: The set of targets to draw that were used in Study 1: a) Circle. b) Line. c) Bow. d) Tilde. e) Ellipse.

4.2.3 Procedures. We aim to gauge the efficacy of the tool through a triad of metrics: stroke precision, stroke quality, and efficiency.

Stroke precision is defined by the degree to which a user's stroke adheres to the intended task outcome. We measured the error, defined as the average distance from the point drawn by the user to the closest point on the target stroke. For circle or horizontal line tasks, it can be simplified to the distance from the point to the axis origin or X-axis. As for other tasks, it was approximated by the distance of the nearest discrete point on the target strokes. Stroke quality typically encompasses attributes such as the smoothness and aesthetic appeal of the lines. Within this context, we quantify stroke

quality by the degree of non-smoothness which was calculated as the mean difference of curvature of the drawn curve stroke (exclude line task). a lower measure of non-smoothness indicates a higher quality, characterized by smoother strokes and reduced jaggedness. Jaggedness is usually caused by the vibration of the brush controller during use. In terms of efficiency, we measured the completion time.

4.2.4 Participants. We recruited paid 16 participants (eight female and eight males, avg. age = 24.2 years) from a local university, and all of them is right-handed. Among them, five participants had no or little VR experience. None of the participants had prior sketch experience in VR.

Participants signed a consent form before starting the study. They were given a tutorial on how to use the default drawing method provided by the using WieldingCanvas, OnlyCanvas and TiltBrush respectively. Then, they were given a few trials as training before doing the tasks. In the experiments, the participants were asked to complete all task using both WieldingCanvas and TiltBrush. The order of the techniques was counter-balanced. It took an average of 50 minutes for each participant to complete the tasks including the training phase. Besides the error and smoothness, we measured the task completion time for all task. The completion time of every task was measured as from when the target appeared, to when the user clicked the submit button. To ensure experimental fairness, all condition were conducted using the same type of brush and color, sourced from the open-source version of Tilt Brush. Additionally, the brush thickness was set to a consistent value across condition.

4.2.5 Results. The statistical significance of the findings was evaluated using a Friedman's test, followed by pairwise comparisons using Bonferroni Correction.

Error bias was statistically significantly different using different 3 methods, $\chi^2(2) = 204, p < 0.01$. Post-hoc pairwise comparison showed significant differences. The deviation of WeildingCanvas ($M = 0.56\text{cm}$, $SD = 0.49\text{cm}$) is significantly lower than OnlyCanvas ($M = 0.66\text{cm}$, $SD = 0.43\text{cm}$), $p < 0.05$, and significantly lower than TiltBrush ($M = 1.61\text{cm}$, $SD = 1.28\text{cm}$), $p < 0.01$. OnlyCanvas is also significantly lower than TiltBrush , $p < 0.01$.

Friedman's test showed a statistically significant effect on curve task (exclude "line" task) on different 3 methods of non-smoothness, $\chi^2(2) = 276.58, p < 0.01$. Post-hoc pairwise comparison showed significant differences. The WeildingCanvas ($M = 0.27$, $SD = 0.19$) is significantly smoother than OnlyCanvas ($M = 0.71$, $SD = 0.10$), $p < 0.01$, and significantly smoother than TiltBrush ($M = 0.52$, $SD = 0.08$), $p < 0.01$. TiltBrush is significantly smoother than OnlyCanvas , $p < 0.01$.

Friedman's test was conducted to determine whether there was a statistically significant effect of drawing methods on completion time. completion time was statistically significantly different using different 3 methods, $\chi^2(2) = 131.43, p < 0.01$. WieldingCanvas ($M = 16.93\text{s}$, $SD = 12.35\text{s}$) took longer than both OnlyCanvas ($M = 9.16\text{s}$, $SD = 5.98\text{s}$), $p < 0.01$ and TiltBrush ($M = 7.73\text{s}$, $SD = 5.44\text{s}$), $p < 0.01$. TiltBrush is significantly faster than OnlyCanvas , $p < 0.01$. The more details of each task can be seen in the Figure 15.

In our analysis of Study 1, we observed that the low levels of non-smoothness and error indicators strongly support our initial hypothesis. These results indicate a high degree of control and precision in the participants' strokes, which aligns with our predicted




Figure 15: Objective evaluation index in Study 1.

outcomes. The low non-smoothness scores suggest that the strokes made by the participants were consistent and high quality, a key aspect of our hypothesis that posits improved drawing ability with the use of our tool. Similarly, the minimal error rates further corroborate this, demonstrating a close alignment between the intended and actual outcomes of the drawing tasks. WeildingCanvas and OnlyCanvas demonstrate lower error bias compared to TiltBrush, which is not surprising given that both condition utilize virtual drawing boards to mitigate the impact of depth. In WeildingCanvas, strokes are primarily built through the movement of the canvas itself, allowing the stylus to remain stationary in mid-air. By allowing the stylus to remain static, WeildingCanvas eliminates the need for users to maintain a precise distance between the stylus and the virtual canvas, reducing the likelihood of jitteriness and enhancing the overall smoothness and precision of the drawing.

Interestingly, while OnlyCanvas has a lower error bias compared to TiltBrush, it falls short in terms of smoothness, as reported by users. One particular evaluation by a user, identified as P4, noted, "When I was using OnlyCanvas in an attempt to accurately trace target strokes, I felt a noticeable jitter." This subjective feedback aligns with the statistical results. Indeed, while OnlyCanvas has successfully mitigated the impact of depth through its virtual drawing board, the absence of physical support appears to pose a challenge for users in maintaining consistent distance between the stylus and the virtual canvas. This lack of physical reference contributes to subtle jitteriness during the brush moving process, adversely affecting the smoothness of the strokes.

The extended completion time on WeildingCanvas could be attributed to several factors. For instance, the method WeildingCanvas employs, where strokes are constructed through the movement of the canvas while the stylus remains stationary, might be more time-consuming, despite its benefits in precision and smoothness. On the other hand, OnlyCanvas and TiltBrush might allow for faster interactions, possibly due to methods of drawing. The fact that OnlyCanvas minimizes depth effects might allow users to draw more quickly without being overly concerned about precision.

OnlyCanvas occupies a middle-ground position among the three platforms across the metrics of error bias, and completion time. The middling error bias and precision levels in OnlyCanvas might be attributable to its virtual drawing board, which eliminates depth effects but lacks physical support, causing some imprecision and jitter. Such a balanced profile might make OnlyCanvas a versatile

option, appealing to users who want a compromise between precision and speed. While it doesn't excel in any particular metric like WeildingCanvas does in precision and smoothness, nor does it facilitate faster task completion like TiltBrush, it offers a middle-of-the-road alternative that could suit a wide range of users with diverse needs and preferences.

4.3 Study 2: Evaluating Sketching Tools

In this study, users had the opportunity to utilize WeildingCanvas to freely create and stitch together canvases within a virtual space. The performance was then compared with sketches made using TiltBrush. The primary focus of this study was to explore the user experience of constructing three-dimensional objects with WeildingCanvas in VR. Based on the findings from Study 1, we decided not to include "OnlyCanvas" as an experimental condition in the current study.

4.3.1 Participants. We recruited paid 16 participants (seven female and nine males, avg. age = 23.4 years) from a local university, and one of them is left-handed and the others are right-handed. Among them, nine participants had no or little VR experience, and two participants had proficient VR experience.

4.3.2 Tasks. In Study 2, participants were asked to do freehand sketching of curves of different levels of complexity. Participants were shown the target pictures and were asked to draw them using WeildingCanvas and TiltBrush, following a methodology akin to that of Barrera Machuca et al.[25, 27]. This approach is more beneficial than simply having participants recreate 3D objects, as it not only allows for observing their capability to independently use the canvas properties to construct 3D drawings but also provides insights into their sketching planning process. We chose three commonly seen drawing examples of different categories of daily objects with moderate complexity appropriate for the completion time of the study. The objects consist of combinations of complex curvatures and lines, representative of many common drawing targets. Figure 16 shows the target objects, which have a clear and easily observable structure.

4.3.3 Procedure. Before participating in the study, all individuals completed a consent form. An initial tutorial acquainted them with the standard drawing functionalities of both TiltBrush and WeildingCanvas. Participants then went through several practice rounds prior to embarking on the formal tasks. The experiment required




Figure 16: Study 2 investigated three different objects. a) Book. b) Chair. c) Sapling.

that they accomplish all tasks utilizing both WieldingCanvas and TiltBrush, with the sequence of technology alternated to maintain balance. There was a five-minute intermission between each technology, and the entire procedure, including training, averaged about one hour per participant. In the second part of the study, timing commenced when participants signaled their readiness by pressing any key and concluded when they actively pressed the 'submit' key down. For left-handed users, we will swap the functions between the left and right controllers, allowing users to draw with their dominant hand.

After completing the tasks, the participants were asked to complete both NASA TLX and rate the following statements on a 5-point Likert scale (1: strongly disagree; 5: strongly agree): 1) *The method was intuitive to use*; 2) *The system helped me express my intended shapes correctly*; 3) *The system helped me draw shapes accurately*; 4) *The drawing functions were easy to use*; 5) *The drawing functions were easy to control*. The participants were asked to evaluate the UI design as well. Finally, they were interviewed about their experiences of using WieldingCanvas, and asked to provide further feedback about their ratings. The interview focused on their usage patterns and expectations for further improvement.

4.3.4 Results.

We analyzed the completion time of the tasks using WieldingCanvas and TiltBrush (Figure 17). In study 2, WieldingCanvas ($M=3.51\text{min}$, $SD=1.54\text{min}$) took longer than TiltBrush ($M=2.19\text{min}$, $SD=1.23\text{min}$). The tests showed a significant effect of WieldingCanvas on completion time for study 2 ($t=2.673$, $p<0.05$). This outcome was expected, as WieldingCanvas required participants to perform additional steps to invoke the canvases before drawing. In contrast, the TiltBrush option allowed them to sketch directly, thus reducing the completion time.

Despite the additional cost of time, WieldingCanvas received positive feedback in general from participants, who appreciated its additional functionality and flexibility.

User generated sketches. Our study data, as illustrated in Figure 18, corroborates our hypothesis that WieldingCanvas can enhance the quality of complex 3D hand-drawn objects. It can be seen that using WieldingCanvas, the participants were able to draw more smooth and accurate lines and curves in VR, which were reflected in the details of the drawings.

Likert-scale ratings. There was no significant improvement ($p>0.05$) observed in using canvas samples for the rating of MD (Mental Demand), FD (Physical Demand), TD (Temporal Demand),




Figure 17: Objective evaluation index in Study 2.

and Ef (Effort) dimensions in the TLX scale. Using WieldingCanvas resulted in significant difference for Pe (Performance) ($t=2.677$, $p=0.012$) and Fr (Frustration) ($t=-2.197$, $p=0.036$), and The comparative analysis reveals that WieldingCanvas's average Pe ($M=15.88$, $SD=3.76$) is significantly higher than TiltBrush's ($M=11.69$, $SD=5.0$), while WieldingCanvas's Fr ($M=3.88$, $SD=2.80$) is significantly lower than that of TiltBrush ($M=6.69$, $SD=4.29$), as shown in Figure 19.

The results indicate that the use of the WieldingCanvas sample is not significant ($p>0.05$) for Q1, Q4, and Q5, suggesting that the sketchpad sample is consistent for these questions and there is no discernible difference. On the other hand, using the drawing board for Q2 ($t=5.076$, $p<0.01$) and Q3 ($t=5.679$, $p<0.01$) yields a significant difference at the 0.01 level. Specifically, the ratings of WieldingCanvas for Q2 ($M=4.63$, $SD=0.62$) and Q3 ($M=4.38$, $SD=0.62$) are significantly higher than the ratings for TiltBrush for Q2 ($M=2.75$, $SD=1.34$) and Q3 ($M=2.69$, $SD=1.01$).

Qualitative feedback. We collected qualitative feedback through semi-structured interviews at the end of the study. Integrating questionnaire feedback and scores, we confirmed the hypothesis that WieldingCanvas is an effective tool. Interviews show that participants generally like the idea of using a virtual canvas to paint in VR. A sample comment given by a participant: "painting with the help of a virtual canvas is a great idea and it helps me control my strokes very well" (P5). Regarding the WieldingCanvas condition, P6 commented "In actual use, I can better draw the shape I want with WieldingCanvas". P1 who is good at oil painting said: "WieldingCanvas allows me to quickly transfer the experience of traditional painting to VR painting". However, another participant expressed that "some of the curves I prefer to paint directly on the canvas, although the effect is not as good as bending canvas" (P2). Regarding the TiltBrush condition, the participants agreed that "the trajectory in the air is difficult to control" (P8, P10, P11).




Figure 18: Example sketches created in Study 2 using both WeldingCanvas and TiltBrush .

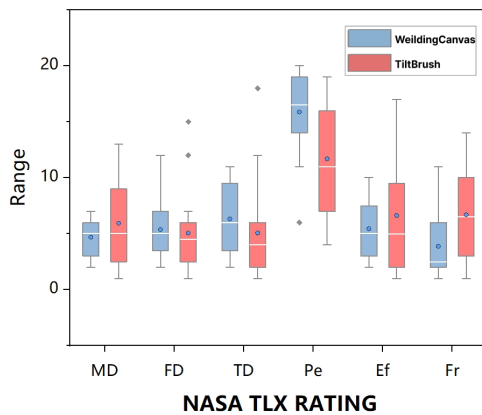


Figure 19: TLX Rating for the overall task.

Through the questionnaire, the participants' qualitative feedback was in general positive, with users preferring WeldingCanvas for its being easy and intuitive to use. The participants answered unanimously that they strongly agree that "the WeldingCanvas helped them draw what they want accurately" with Q2. The sentiment was positive-to-neutral for all questions in Q2. P9, who did not have previous VR experience, commented, "The drawing with WeldingCanvas was very interesting." P12 remarked, "I am satisfied with the drawing results by using WeldingCanvas." P7 said, "Although it helped drawing with WeldingCanvas, it took more time."

Most participants unanimously agree that "it is easy to draw by WeldingCanvas" (Q5). According to the results of the interview, participants all prefer to draw by WeldingCanvas. while some users have expressed a desire for the capability to draw directly on the canvas. P7 said, "For tasks that don't demand precise line drawing, I'd rather draw directly on the canvas than rely on moving the canvas to create lines. Subsequently, I can also make use of the canvas transformation features for further editing. This method saves time throughout my creative process."

Q1 asked "whether the method is intuitive", and most of the participants agreed with it, and some participants maintained a

neutral attitude. P9 remarked that "Folding of the drawing canvas is intuitive and easy to use". P3 believed that "the different canvas can be used as references so that I can intuitively estimate the spatial position". While P7 thought that "sometimes the drawing effect can't be seen right away, and I need additional operations".

4.4 Sketching Activity: Exploring Freeform Sketching with WeldingCanvas

In this section, we initiated a specially sketching activity themed 'Exploring Freeform Sketching with WeldingCanvas'. This sketching activity imposed no restrictions, allowing participants to fully exercise their creative potential. The purpose of this approach was to gain a more accurate understanding of how users interact with the drawing tools and features when given complete autonomy, thereby providing a more natural and realistic assessment of the platform's capabilities.

We recruited 8 paid participants to complete the drawing creation (four female and three males, avg. age = 24.2 years), two of whom have extensive drawing experience. After an hour of training, each participant was provided with a Meta Quest 2 equipped with the WeldCanvas app. Over a span of three days, participants had the flexibility to schedule their own time to fully engage with our system and were required to record their process while using it. We collected both the artworks created by the users and videos of their drawing processes.

As shown in Figure 21, we have selected a few works that are particularly distinctive. These drawings were done on different canvases and were stitched together using the canvas management feature. These works made minimal use of the built-in drawing methods provided by Tilt Brush.

The Chinese knot shown in Figure 21 (a) was created by P1, a novice in drawing. The main body and grid part of the Chinese knot were made using the Line feature, accomplished through the movement of the canvas and minor adjustments by the artist. P1 commented, "The drawing functions of WeldingCanvas were a very pleasant surprise for me. They helped me create a complete piece of art". In addition to that, most parts of the Chinese knot were drawn using WeldingCanvas's canvas feature. The connecting parts were created using the symmetry function, and the ribbons were crafted using arcs combined with transformations. Except for the tassel at

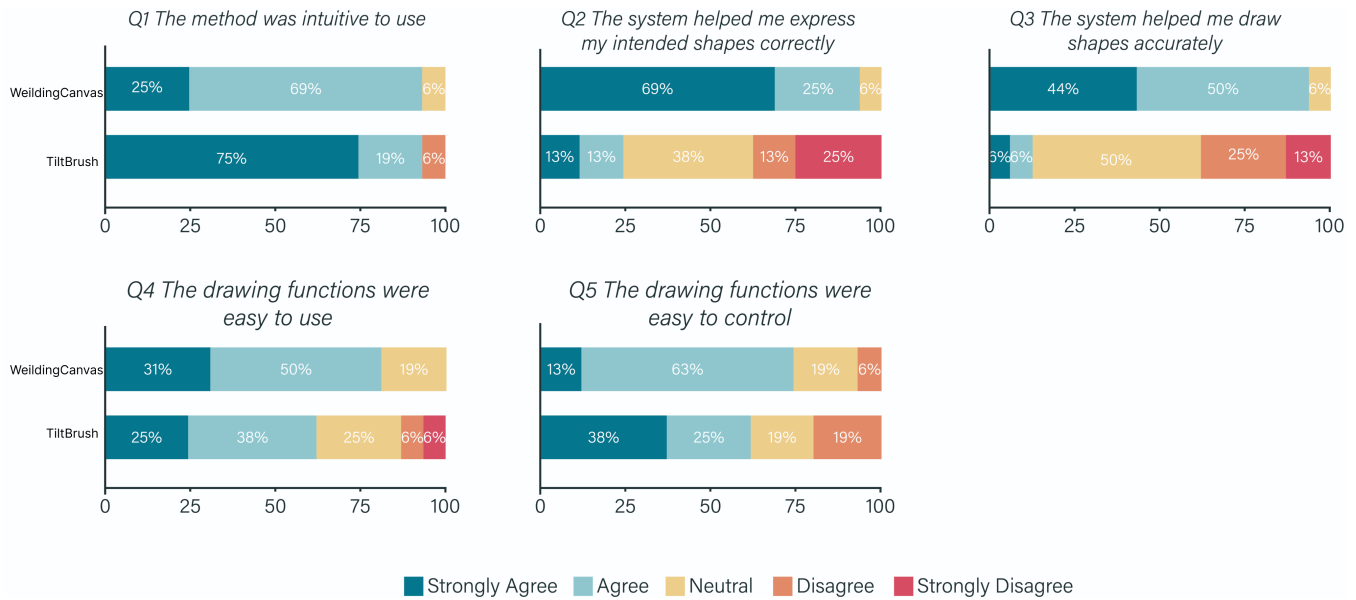


Figure 20: The answers to the questionnaire using a 5-point Likert scale. Red shades indicate disagreement, orange is neutral and green shades represent agreement.




Figure 21: Example drawings created in Sketching Activity.

the bottom and the text portion, which were drawn using multiple OnlyCanvas and then stitched together. The artworks displayed in Figure 21 (b-d) were created by participants with prior drawing experience. The majority of the leaves and flowers shown were drawn on OnlyCanvas. These elements were then transformed into three-dimensional curved structures using WieldingCanvas’s transformation features.

Finally, they were stitched together using the canvas management function. All participants unanimously agreed that this method of drawing on a virtual canvas, combined with canvas transformations or canvas stitching, offers a unique blend of two-dimensional drawing and three-dimensional interaction. They also concurred

that achieving these kinds of artworks would have been difficult using previous applications.

5 DISCUSSION

WieldingCanvas was presented with the primary goal of creating 3D drawing lines and curves easily and accurately. We investigated canvas manipulations to assist the drawing of lines and curves in 3D. A user study was conducted to evaluate the capability and usability of WieldingCanvas. Our study indicates that WieldingCanvas has a significant positive impact on drawing smoothness and reducing errors. Most participants preferred drawing with WieldingCanvas,

though it spent more time to complete the drawings compared with directly sketching in VR.

Compared with TiltBrush, WeldingCanvas resulted in better smoothness and precision. The virtual plane of WeldingCanvas provides a useful spatial reference for the user during drawing, mitigating issues related to depth perception errors to a certain extent. In comparison to OnlyCanvas, the moving canvas feature of WeldingCanvas contributed to a noticeable improvement in both precision and smoothness.

As mentioned, WeldingCanvas may have taken longer to complete tasks because it offered users more control and flexibility, allowing them to adjust the canvas and draw with more precision. In contrast, TiltBrush was more simplistic and allowed for more casual sketching. Therefore, users may have been able to complete tasks more quickly with TiltBrush because they did not have to worry about adjusting the canvas or drawing with as much precision.

This would also ask the novel users to spend more time to learn the skills. Besides, asking users to draw on the canvas itself adds a certain restriction that could increase the time. Based on the performance of the expert users, it can be seen that users can use WeldingCanvas with higher efficiency after sufficient training. On the other hand, the taskload ratings based on the participants were lower than TiltBrush, regardless of the time completion time. We think it is caused by users' habit of writing and drawing on the plane, and such life experiences makes the tasks less stressful. In the questionnaire, Q2, Q3 were rated higher than TiltBrush but Q1, Q4, Q5 were rated lower. These also prove the previous results, that the improved sketching smoothness and precision helped users better create the drawings.

No participant in our study reported disturbances in their sketching process when using WeldingCanvas. However, we did observe that some planning is required in coordinating canvas motion with drawing actions. This aspect, along with the trial-and-error process inherent in mastering any new tool, does introduce a learning effort. Users might need time to adapt to this integrated approach of planning and drawing. While this learning phase might initially impact the fluidity of the creative process, its long-term effect on creativity and workflow efficiency remains an open question. Future studies are needed to explore whether, and to what extent, this learning curve and planning requirement might interfere with users' creative processes over time. The balance between the intuitive, familiar aspects of the tool and the new skills it requires is a crucial area for further investigation.

Our main target in designing WeldingCanvas was to create an accurate 3D sketch with the help of a virtual canvas that leverages the actions of both hands. The poor depth perception in VR is one of the main reasons for the difficulty of accurate sketching. With WeldingCanvas in VR, the user's experience with 2D sketching can be directly migrated to 3D drawing. In the experiments, to compare with sketching in the air we do not support the ability of automatic calibration. As a result, most users adapted directly to drawing on the virtual drawing canvas and created 3D curves that were more accurate than drawing in the air.

The proposed canvas manipulations can be further expanded, as we find the current ones are not suitable for all curves. Through observing users' creation activities in the study, we found that when

drawing shapes like lines and circles, users tended to use the canvas function of WeldingCanvas, while for some other curves, such as irregular arcs, users tended to draw directly on the canvas. In addition, the perception of using WeldingCanvas is very different from typical freehand sketching. When users use WeldingCanvas to create complex shapes, it feels more like splicing multiple clip-arts together in space, and The Snapping Two Canvases function makes this process more convenient with added spatial constraints. Indeed, incorporating snapping functions risks a CAD-like experience. However, to meet certain user needs and enhance usability, their inclusion was necessary. Users have the option to enable or disable these functions. We are considering providing an intermediate state that melds sketch-like and CAD-like features

WeldingCanvas, while having certain limitations, serves more as a complement to existing sketch tools rather than a replacement. The effective synergy between WeldingCanvas and TiltBrush, as demonstrated in our Sketching Activity, supports this viewpoint. There are exciting opportunities for continuing to improve WeldingCanvas for sketching in 3D. Besides the lines and curves, WeldingCanvas can be used in industrial design processes, and allows users to create complex 2D graphics on a virtual reference plane and assemble them spatially.

6 CONCLUSION

The recent advancement of Virtual Reality (VR) technology gives rise to more opportunities that enable users to create 3D shapes. 3D sketching is an intricate topic with the challenge of the absence of support and virtual depth control. This paper introduces WeldingCanvas, an interactive VR sketching canvas that leverages users' spatial manipulation of the canvas to assist 3D sketching. Our work contributes to the VR community with the design and implementation of a new approach that makes 3D sketching more accurate and smooth. In compliance with the open-source agreement of Tilt Brush, we plan to make WeldingCanvas available for app download on the Meta Quest platform in the future.

This paper contributes the design of WeldingCanvas, interaction techniques, and empirical data on the precision of drawing as an extensive exploration of the canvas. WeldingCanvas is shown to be effective in assisting users for accurate 3D sketching in VR, resulting in improved drawing precision and smoothness. Even though it sacrificed some efficiency compared to freehand sketching, the capabilities and interactions of WeldingCanvas made it preferred by the participants.

ACKNOWLEDGMENTS

This work was supported by the National Key R&D Program of China (Grant No. 2022ZD0118002), the National Natural Science Foundation of China (Grant No. 62132010, 62272324), the Municipal Natural Science Foundation of Beijing of China (Grant No. 4222023) and Key Laboratory of Smart Education of Guangdong Higher Education Institutes, Jinan University (Grant No. 2022LSYS003).

REFERENCES

- [1] 2020. Tilt Brush README. <https://github.com/googlevr/tilt-brush> original-date: 2020-04-13T18:55:08Z.
- [2] 2023. The Fascinating Mechanics of a Pottery Wheel – Artabys. <https://artabys.com/the-fascinating-mechanics-of-a-pottery-wheel/>

- [3] Rahul Arora. 2020. Creative Expression with Immersive 3D Interactions. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [4] 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*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3173574.3173759>
- [5] Rahul Arora, Rubaiat Habib Kazi, Fraser Anderson, Tovi Grossman, Karan Singh, and George Fitzmaurice. 2017. Experimental Evaluation of Sketching on Surfaces in VR. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 5643–5654. <https://doi.org/10.1145/3025453.3025474>
- [6] Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2008. iLoveSketch: As-Natural-as-Possible Sketching System for Creating 3d Curve Models. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (Monterey, CA, USA) (UIST '08). Association for Computing Machinery, New York, NY, USA, 151–160. <https://doi.org/10.1145/1449715.1449740>
- [7] Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2009. EverybodyLovesSketch: 3D Sketching for a Broader Audience. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology* (Victoria, BC, Canada) (UIST '09). Association for Computing Machinery, New York, NY, USA, 59–68. <https://doi.org/10.1145/1622176.1622189>
- [8] Ravin Balakrishnan, George Fitzmaurice, Gordon Kurtenbach, and William Buxton. 1999. Digital Tape Drawing. In *Proceedings of the 12th Annual ACM Symposium on User Interface Software and Technology* (Asheville, North Carolina, USA) (UIST '99). Association for Computing Machinery, New York, NY, USA, 161–169. <https://doi.org/10.1145/320719.322598>
- [9] Mayra Donaji Barrera Machuca, Wolfgang Stuerzlinger, and Paul Asente. 2019. The Effect of Spatial Ability on Immersive 3D Drawing. In *Proceedings of the 2019 on Creativity and Cognition* (San Diego, CA, USA). Association for Computing Machinery, New York, NY, USA, 173–186. <https://doi.org/10.1145/3325480.3325489>
- [10] Joachim Deisinger, Roland Blach, Gerold Wesche, Ralf Breining, and Andreas Simon. 2000. Towards Immersive Modeling - Challenges and Recommendations: A Workshop Analyzing the Needs of Designers. In *Proceedings of the 6th Eurographics Conference on Virtual Environments* (Amsterdam, The Netherlands) (EG VE'00). Eurographics Association, Goslar, DEU, 145–156.
- [11] Tobias Drey, Jan Gugenheimer, Julian Karlbauer, Maximilian Milo, and Enrico Rukzio. 2020. *VRSketchIn: Exploring the Design Space of Pen and Tablet Interaction for 3D Sketching in Virtual Reality*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376628>
- [12] Mark Fuge, Mehmet Ersin Yumer, Gunay Orbay, and Levent Burak Kara. 2012. Conceptual design and modification of freeform surfaces using dual shape representations in augmented reality environments. *Computer-Aided Design* 44, 10 (2012), 1020–1032. <https://doi.org/10.1016/j.cad.2011.05.009> Fundamentals of Next Generation CAD/E Systems.
- [13] Richard Fung, Edward Lank, Michael Terry, and Celine Latulipe. 2008. Kinematic Templates: End-User Tools for Content-Relative Cursor Manipulations. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (Monterey, CA, USA) (UIST '08). Association for Computing Machinery, New York, NY, USA, 47–56. <https://doi.org/10.1145/1449715.1449725>
- [14] Google. 2022. *Tilt Brush*. Retrieved March 29, 2022 from <https://www.tiltbrush.com/>
- [15] Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. 2002. Creating Principal 3D Curves with Digital Tape Drawing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Minneapolis, Minnesota, USA) (CHI '02). Association for Computing Machinery, New York, NY, USA, 121–128. <https://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. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) (CHI '03). Association for Computing Machinery, New York, NY, USA, 185–192. <https://doi.org/10.1145/642611.642645>
- [17] Rorik Henrikson, Bruno Araujo, Fanny Chevalier, Karan Singh, and Ravin Balakrishnan. 2016. Multi-Device Storyboards for Cinematic Narratives in VR. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 787–796. <https://doi.org/10.1145/2984511.2984539>
- [18] Rorik Henrikson, Bruno De Araujo, Fanny Chevalier, Karan Singh, and Ravin Balakrishnan. 2016. Stereoboard: Sketching Stereoscopic Storyboards. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 4587–4598. <https://doi.org/10.1145/2858036.2858079>
- [19] Riwano Ikeda and Issei Fujishiro. 2021. SpiCa: Stereoscopic Effect Design with 3D Pottery Wheel-type Transparent Canvases. In *SIGGRAPH Asia 2021 Technical Communications*. 1–4.
- [20] 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 (2007), 1067–1081. <https://doi.org/10.1109/TVCG.2007.1060>
- [21] Daniel F. Keefe, Daniel Acevedo Feliz, Tomer Moscovich, David H. Laidlaw, and Joseph J. LaViola. 2001. CavePainting: A Fully Immersive 3D Artistic Medium and Interactive Experience. In *Proceedings of the 2001 Symposium on Interactive 3D Graphics (I3D '01)*. Association for Computing Machinery, New York, NY, USA, 85–93. <https://doi.org/10.1145/364338.364370>
- [22] Yongkwan Kim, Sang-Gyun An, Joon Hyub Lee, and Seok-Hyung Bae. 2018. *Agile 3D Sketching with Air Scaffolding*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3173812>
- [23] Yongkwan Kim and Seok-Hyung Bae. 2016. SketchingWithHands: 3D Sketching Handheld Products with First-Person Hand Posture. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 797–808. <https://doi.org/10.1145/2984511.2984567>
- [24] Mayra Donaji Barrera Machuca, Rahul Arora, Philipp Wacker, Daniel Keefe, and Johann Habakuk Israel. 2023. Interaction Devices and Techniques for 3D Sketching. In *Interactive Sketch-based Interfaces and Modelling for Design*. River Publishers, 195–239.
- [25] Mayra D. Barrera Machuca, Paul Asente, Wolfgang Stuerzlinger, Jingwan Lu, and Byungmoon Kim. 2018. Multiplanes: Assisted Freehand VR Sketching. In *Proceedings of the Symposium on Spatial User Interaction* (Berlin, Germany) (SUI '18). Association for Computing Machinery, New York, NY, USA, 36–47. <https://doi.org/10.1145/3267782.3267786>
- [26] Mayra Donaji Barrera Machuca, Johann Habakuk Israel, Daniel F. Keefe, and Wolfgang Stuerzlinger. 2023. Toward More Comprehensive Evaluations of 3D Immersive Sketching, Drawing, and Painting. *IEEE Transactions on Visualization and Computer Graphics* (2023), 1–18. <https://doi.org/10.1109/TVCG.2023.3276291>
- [27] Mayra D. Barrera Machuca, Wolfgang Stuerzlinger, and Paul Asente. 2019. Smart3DGuides: Making Unconstrained Immersive 3D Drawing More Accurate. In *25th ACM Symposium on Virtual Reality Software and Technology* (Parramatta, NSW, Australia) (VRST '19). Association for Computing Machinery, New York, NY, USA, Article 37, 13 pages. <https://doi.org/10.1145/3359996.3364254>
- [28] Bearfoot OT. 2021. *What is bilateral coordination and why does it matter?* Retrieved March 29, 2022 from <https://www.bearfootoccupationaltherapy.com/post/what-is-bilateral-coordination-and-why-does-it-matter>
- [29] Ken Pfeuffer, Abdullatif Dinc, Jan Obernolte, Radiah Rivu, Yasmeen Abdrabou, Franziska Shelter, Yomna Abdelrahman, and Florian Alt. 2021. Bi-3D: Bi-Manual Pen-and-Touch Interaction for 3D Manipulation on Tablets. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '21). Association for Computing Machinery, New York, NY, USA, 149–161. <https://doi.org/10.1145/3472749.3474741>
- [30] Steven Schkolne, Michael Pruett, and Peter Schröder. 2001. Surface Drawing: Creating Organic 3D Shapes with the Hand and Tangible Tools. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, Washington, USA) (CHI '01). Association for Computing Machinery, New York, NY, USA, 261–268. <https://doi.org/10.1145/365024.365114>
- [31] Johannes Schmid, Martin Sebastian Senn, Markus Gross, and Robert W. Sumner. 2011. OverCoat: An Implicit Canvas for 3D Painting. In *ACM SIGGRAPH 2011 Papers* (Vancouver, British Columbia, Canada) (SIGGRAPH '11). Association for Computing Machinery, New York, NY, USA, Article 28, 10 pages. <https://doi.org/10.1145/1964921.1964923>
- [32] Karan Singh. 1999. Interactive Curve Design Using Digital French Curves. In *Proceedings of the 1999 Symposium on Interactive 3D Graphics* (Atlanta, Georgia, USA) (I3D '99). Association for Computing Machinery, New York, NY, USA, 23–30. <https://doi.org/10.1145/300523.300525>
- [33] Gravity Sketch. 2022. *Gravity Sketch*. Retrieved March 29, 2022 from <https://www.gravitysketch.com/>
- [34] Floor Verhoeven and Olga Sorkine-Hornung. 2019. Rodmesh: Two-handed 3D surface modeling in virtual reality. (2019).
- [35] Philipp Wacker, Adrian Wagner, Simon Voelker, and Jan Borchers. 2018. Physical Guides: An Analysis of 3D Sketching Performance on Physical Objects in Augmented Reality. In *Proceedings of the Symposium on Spatial User Interaction* (Berlin, Germany) (SUI '18). Association for Computing Machinery, New York, NY, USA, 25–35. <https://doi.org/10.1145/3267782.3267788>
- [36] E. Wiese, J. H. Israel, A. Meyer, and S. Bongartz. 2010. Investigating the Learnability of Immersive Free-Hand Sketching. In *Proceedings of the Seventh Sketch-Based Interfaces and Modeling Symposium* (Annecy, France) (SBIM '10). Eurographics Association, Goslar, DEU, 135–142.
- [37] Min Xin, Ehud Sharlin, and Mario Costa Sousa. 2008. Napkin Sketch: Handheld Mixed Reality 3D Sketching. In *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology* (Bordeaux, France) (VRST '08). Association for Computing Machinery, New York, NY, USA, 223–226. <https://doi.org/10.1145/1450579.1450627>
- [38] Emilie Yu, Rahul Arora, Tibor Stanko, J. Andreas Børentzen, Karan Singh, and Adrien Bousseau. 2021. CASSIE: Curve and Surface Sketching in Immersive Environments. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 190, 14 pages. <https://doi.org/10.1145/3411734.3411744>

3411764.3445158

- [39] Xue Yu, Stephen DiVerdi, Akshay Sharma, and Yotam Gingold. 2021. ScaffoldSketch: Accurate Industrial Design Drawing in VR. In *The 34th Annual*

ACM Symposium on User Interface Software and Technology (Virtual Event, USA) (UIST '21). Association for Computing Machinery, New York, NY, USA, 372–384. <https://doi.org/10.1145/3472749.3474756>