Agile 3D Sketching with Air Scaffolding

Yongkwan Kim, Sang-Gyun An, Joon Hyub Lee, Seok-Hyung Bae
Department of Industrial Design, KAIST
yongkwan.kim | sang-gyun.an | joonhyub.lee | seokhyung.bae @ kaist.ac.kr

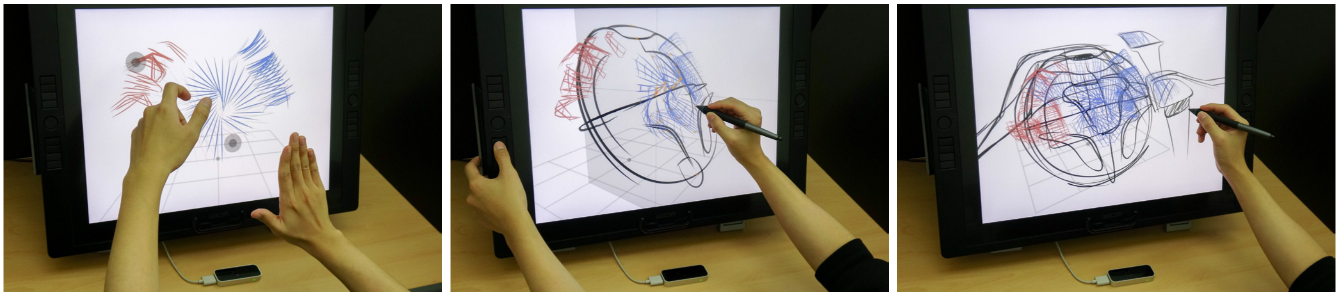

(a) (b) (c)

Figure 1. In our agile 3D sketching workflow with air scaffolding, in which hand motion and pen drawing complement each other, (a) the user makes unconstrained hand movements in the air to quickly generate rough shapes to be used as scaffolds, (b) uses the scaffolds as references and draws finer details with them (c) in an iterative and progressive manner to produce a high-fidelity 3D concept sketch of a steering wheel.

ABSTRACT

Hand motion and pen drawing can be intuitive and expressive inputs for professional digital 3D authoring. However, their inherent limitations have hampered wider adoption. 3D sketching using hand motion is rapid but rough, and 3D sketching using pen drawing is delicate but tedious. Our new 3D sketching workflow combines these two in a complementary manner. The user makes quick hand motions in the air to generate approximate 3D shapes, and uses them as scaffolds on which to add details via pen-based 3D sketching on a tablet device. Our air scaffolding technique and corresponding algorithm extract only the intended shapes from unconstrained hand motions. Then, the user sketches 3D ideas by defining sketching planes on these scaffolds while appending new scaffolds, as needed. A user study shows that our progressive and iterative workflow enables more agile 3D sketching compared to ones using either hand motion or pen drawing alone.

Author Keywords

Hand motion; scaffolding; 3D sketching; product design

ACM Classification Keywords

H.5.2. Information interfaces and presentation (e.g., HCI): User Interfaces-Input devices and strategies, Interaction styles; I.3.8. Computer Graphics: Applications

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

ACM ISBN 978-1-4503-5620-6/18/04…$15.00
https://doi.org/10.1145/3173574.3173812

INTRODUCTION

Hand motion is an intuitive and useful mode of expression [30]. People adroitly make hand motions freely in the air to communicate the scale and features of 3D shapes [10]. Such expressiveness can be leveraged to help designers in the early stages of rapid and rough idea exploration, for which conventional 3D CAD modeling systems are unsuitable [8, 22, 28]. While 3D sketching using hand motion can help form the approximate shape and scale of a product, it may not meet the level of fidelity required in conceptualization [22]. As such, designers cannot solely rely on their hand motions to develop a fully-fledged design concept.

The pen is another intuitive and expressive tool [20]. Designers trained in traditional drawing techniques can effectively describe 3D shapes with a relatively small number of curves. Thus, there have been attempts to transfer these skills to a digital medium for authoring 3D shapes [6, 20]. However, these techniques introduce varying degrees of additional complexity and indirection in extracting the intended 3D information from the user's 2D input. Maintaining a sense of proportionality and scale is another inherent difficulty while drawing 3D shapes on a 2D medium, whether analog or digital [24].

We introduce a new 3D sketching workflow where the user employs both hand motion and pen drawing (Figure 1). The user creates *air scaffolds*—intermediate shapes containing meaningful spatial information—using unconstrained hand motion in the air and defines sketch planes on them. Then, finer details are added using pen-based 3D sketching techniques. Rather than creating strokes or surfaces directly with raw hand motions, which are too imprecise for product design, we found a *sweet spot* in which rough but quick hand motions create intermediate 3D references for more precise 3D sketching. Our contributions are as follows. We:

- Propose a 3D sketching workflow combining the strengths of hand and pen input.
- Devise an algorithm to identify descriptive hand motions from transitory and extract air scaffolds from the identified motions.
- Integrate the air scaffolding technique with pen-based 3D sketching into a practical system.
- Evaluate our system with users to verify that air scaffolding facilitates idea exploration through agile 3D sketching.

In the following sections, we summarize existing studies on 3D authoring based on hand motion and pen drawing, propose our air scaffolding algorithm and the integrated 3D sketching workflow, exhibit sample sketches produced using our workflow, explain the implementation in detail, analyze the results from our user study, provide in-depth discussion, and finally make conclusions.

RELATED WORK

To develop a 3D authoring workflow consisting of hand motion and pen sketching that complement each other, we summarize related studies on these two modes of input and identify the strengths and weaknesses of each.

3D Authoring Based on Hand Motion in the Air

Recent developments in hand-tracking technologies [26, 30] have led to various 3D authoring tools that exploit the expressiveness of the human hands [27].

Techniques that directly utilize the posture and motion of bare hands in free space are more effective when coupled with virtual and augmented reality technologies that visualize the 3D paths of the hand motion [8, 15, 22, 29].

In addition, elaborate 3D authoring requires fine motor control that may benefit from haptic feedback [17]. However, many current haptic technologies still severely restrict degrees of freedom and range of movement, undermining the expressiveness of hand motions in space.

Without haptic feedback, users may be unable to create their desired 3D shapes solely based on hand movements in space. Thus, some techniques provide the means to modify motion-generated shapes *post hoc* [8, 22]. However, these can add complexity and require more effort from the user.

While hand motion may be unsuitable for creating high-fidelity 3D shapes without the assistance of appropriate feedback technologies, it may hold important cues regarding the shapes the user had in mind. In the study by Holz and Wilson [10], when people were asked to describe a given 3D shape using bare hand gestures, they made gestures describing the overarching geometry, scale, and proportion of the shape, through various postures and motions of the hands, e.g., with spread palms swept tangentially to depict flat surfaces.

Thus, we see that unconstrained hand motions contain important spatial clues that can be used as a stepping stone toward more elaborate 3D authoring.

3D Authoring Based on Pen Drawing on Tablets

Designers have traditionally used perspective drawings with pen and paper to depict 3D shapes on 2D surfaces. There are ongoing efforts to transfer perspective drawing techniques to create 3D shapes composed of 3D curves through pen input on 2D digital screens [6, 20].

In order to infer 3D shapes from 2D drawings, some techniques require the user to sketch multiple instances of the same curve from different views [2, 14], while others require the user to predefine 3D planes and curved surfaces on which to situate the curves [3, 7, 13, 21, 23, 31]. Many techniques take the latter approach for the simplicity of the mental model and its similarity to the traditional practice.

Many pen-based 3D sketching systems assist the user with discovering and configuring appropriate sketch surfaces through preset sketch planes [7], manual control of the position and orientation of sketch planes using buttons [13] or dedicated 3D input devices [21], pen strokes on existing surfaces to invoke perpendicular sketch planes [31], offsetting and erecting sketch surfaces along a given 3D mesh [5, 23], and bounding boxes often used in architectural drawings [25]. These show that the ease with which the user can define and change the position and orientation of sketch surfaces affects the agility and expressiveness of 3D sketching.

In particular, Bae et al. [3] introduce a tick-based technique in which horizontal, vertical, and oblique sketch planes are defined in reference to existing 3D curves. While this technique enables the user to efficiently define arbitrary sketch planes, one downside is that a reference must first exist for sketch planes to be created relative to it. This is especially problematic during the early phase of sketching, when 3D references are scarce or nonexistent.

SketchingWithHands suggests a remedy to this problem by introducing a virtual hand model captured using a hand-tracking sensor [16]. In this system, a skeletal model of the reconstructed hand is used as a reference against which to define sketch planes. Another benefit is that the hand model provides the needed spatial clues to design handheld objects. While this approach is a novel solution to a particular domain, it lacks general applicability.

We see that pen-based 3D authoring techniques using sketch surfaces offer simplicity and familiarity, and that the ability to quickly define desired sketch planes is crucial. Various techniques show that defining sketch planes in reference to existing 3D entities can be quick and accurate, but a *chicken-or-the-egg* problem arises at the initial stage, when nothing exists to use as the reference. In fact, help is needed whenever the user wishes to make a significant departure from what already exists.

Hybrid Approaches

Recent studies have explored using surficial and spatial inputs complementarily for authoring 3D contents. De Araujo et al. [4] implemented a system in which the user draws a cross-sectional profile on a surface and makes hand motions in space to create a path along which a volume is generated, akin to extrusion in CAD software. Jackson and Keefe [12] devised a system in which paper sketches are suspended in space as references for 3D modeling using hand motion.

Our approach differs in that we use shapes generated using hand motion only as intermediate scaffolds to assist with more precise pen-based 3D sketching, rather than as the final outcome, in consideration of the limited fidelity of the shapes that hand motion can produce.

3D SKETCHING WITH AIR SCAFFOLDING

In this section, we propose an integrated 3D sketching workflow that combines the strengths of the hand motion and pen drawing in a complementary manner.

First, the user makes unconstrained hand motions in the air to generate rough shapes, from which only the intended and descriptive shapes are extracted and processed into workable scaffolds through our algorithm. These scaffolds are then utilized as references for defining 3D sketch planes, on which the user performs 3D perspective drawings to create high fidelity 3D curves. Our progressive and iterative workflow enables the user to add more scaffolds and 3D curves at any stage.

This air scaffolding technique leverages how people describe 3D shapes using hand gestures, as reported by Holz and Wilson [10], in a workflow that organically incorporates state-of-the-art, pen-based 3D sketching techniques.

Defining Hand Profile

In this step, we track all hand movements in real time but isolate intended, descriptive hand motions from transitory hand motions and use them to construct air scaffolds.

Curve network of hand skeleton

We define the basic geometry from raw data captured with a hand-tracking sensor (Figure 2a). A network representing the hand is constructed from four curves running along the lengths of four fingers. We exclude the thumb, as it is usually unused when people make hand movements to describe overall shape and curvature, reserved only for depicting smaller features such as tubes [10].

Determining meaningful motion

Our aim is to automatically identify and isolate the intended, descriptive hand motions apart from transitory hand motions without the user having to explicitly activate and deactivate specific input modes, for unhindered and intuitive interaction.

When asked to describe 3D shapes using hand movements, people tended to pose their palm so that it would fit the curvature of the imaginary shape. They then swept the palm along the imaginary surface [10]. Exploiting this tendency, we determine that the user is performing a descriptive motion only when the palm's movement is tangential to the imaginary surface described by the palm's posture, within a certain tolerance (Figure 2b). Otherwise, we interpret that the user's hand is transiting from one region of interest to another. Through pilot tests, we made a number of heuristic decisions: the relevant vectors were placed at the center of the proximal phalanx of the middle finger, and the maximum tolerance of deviation from the tangent was set to 27.5°.

Hand profile describing shapes

When the hand motion is determined to be descriptive, additional steps are taken to ensure that only the minimum essential data is processed and visualized to avoid clutter and overhead.

We define the generator profile as the intersection between the network of curves representing the postured palm and the plane perpendicular to the velocity vector (acquired from the hand-tracking sensor) and use it to construct an air scaffold (Figure 2c). The extracted generator profile is resampled to contain ten equidistant points along the cord length. Such a resampling ensures that the same number of points are used for hand motion in any direction.

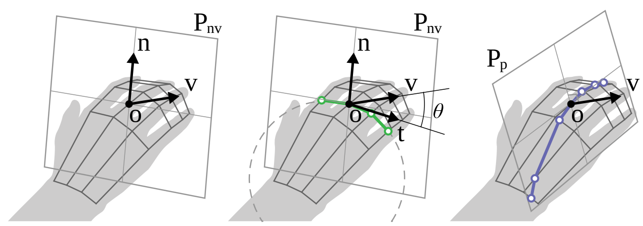

Figure 2. A 4 × 5 curve network is constructed by aggregating four curves that run along four fingers. (a) P_{nv} is a plane that contains both n (up vector) and v (velocity vector) at o (the center point of the middle finger's proximal phalanx). (b) The intersection (green) between P_{nv} and the curve network is fitted to a circle with t (tangent vector) at o. The hand motion is determined to be descriptive if angle θ between v and t is less than a certain threshold, or transitory if it is larger. (c) If the hand motion is descriptive, a generator profile is defined as an intersection (blue) between P_p (a plane passing through o and perpendicular to v) and the curve network.

Air Scaffolding from Hand Gestures

We then use the generator profile to construct an air scaffold. However, interpolating the paths travelled by the dynamically changing generator profile faces two major issues: first, despite our efforts to identify descriptive motion, unintended trails can still be left behind; second, the hand sometimes revisits the imaginary shape and overpopulates a region.

Our approach (Figure 3) solves these two issues simultaneously. We first mark all descriptive trails traveled by the hand as *candidates*. According to Holz and Wilson [10], people tend to make similar gestures repeatedly to describe a shape. Again, we utilize this innate tendency to interpret the user's true intent (Figure 4). Furthermore, the same approach allows the user to modify and adjust the air scaffolds without requiring the user to explicitly specify her intent to do so.

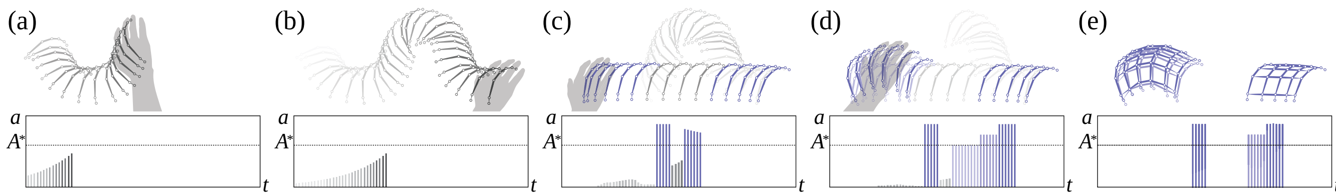

Figure 3. Unconstrained hand motion in free space leaves (a) a trail of polymer particles in the air that (b) decays over time. (c) When the hand repeatedly passes over a region, the quantity of polymer particles exceeds a certain threshold and is thus solidified (blue). (d) As time passes, the solid portion stays constant, whereas the remaining particles continue to decay. When the hand passes near a previously solidified polymer, some parts of the solidified polymer are melted back into polymer particles for re-solidifying. (e) After an extended period of time, all of the polymer particles decay away, leaving only the solidified polymers.

To better illustrate the technical details of our approach, we introduce a metaphor that likens the user's hand to an imaginary 3D printer that sprays, solidifies, and melts diffusive polymer with its heated nozzle in an iterative manner until the user is satisfied with the produced shape.

Spraying available material

The aforementioned generator profile acts as the nozzle of the 3D printer (Figure 3a-b). As the user moves the hand, the nozzle sprays and leaves a trail of polymer particles in the air in the shape of the nozzle. The quantity of particles sprayed by the ith instance of the nozzle a_i decays as a function of time passed since the initial spraying t_i^0:

$$a_i(t; t_i^0) = Ae^{-\lambda(t-t_i^0)} \tag{1}$$

where A is the initial quantity sprayed and λ is the decay constant (we set λ = 0.35 s^{-1}). Here, the half-life is $\ln 2/\lambda$ (~2.0 s).

Solidifying by aggregating the available material

Given the current jth instance of the nozzle, the quantity of polymer particles in the air available for solidifying at the position of the current instance $\tilde{a}_j$ is the aggregate of all available particles sprayed by previous nozzle instances:

$$\tilde{a}_j(t) = \sum_{i=1}^{I} a_i(t; t_i^0) e^{-Bd_{ij}} \tag{2}$$

where d_{ij} is the distance between the ith and jth instances of the nozzle and B is a kind of diffusion constant (we set B = 0.3 cm^{-1}). Here, if the available quantity $\tilde{a}_j$ exceeds a certain threshold A^* (we set A^* = 1.25 A), which specifies the minimum quantity needed for solidification, then the jth instance of the nozzle at the current position will solidify the particles in its shape (Figure 3c). Upon solidification, the quantity of particles transformed into the solid is immediately subtracted from the surroundings. The solidified polymer will no longer decay as a function of time and thus maintain its constant quantity over time. On the other hand, if the available quantity $\tilde{a}_j$ does not exceed A^*, then the nozzle will spray additional particles that will decay in quantity according to equation 1.

Melting solids for re-solidifying

If the nozzle approaches an already solidified polymer, the solid will be partially melted and converted back into decaying polymer particles, which the approaching instance of the nozzle can aggregate (Figure 3d). The closer the nozzle is to the solidified polymer, the more the solid will melt and be converted back to particles. Given the previously solidified quantity A_k, the quantity of particles that is melted from the solid by the current jth nozzle $a_j(t; t_j^0)$ can be expressed as a function of time since their melting:

$$a_j(t; t_j^0) = \frac{1}{d_{kj}^2 + 1} A_k e^{-\lambda(t-t_j^0)} \tag{3}$$

Here, the remaining quantity of the solidified polymer at $t = t_j^0$ is $A'_k = A_k d_{kj}^2 / (d_{kj}^2 + 1)$. Should this quantity fail to exceed the aforementioned threshold A^*, the entire solid will be immediately converted back into available polymer particles, of which the quantity a_{j+1} is expressed as:

$$a_{j+1}(t; t_{j+1}^0) = A'_k e^{-\lambda(t-t_{j+1}^0)} \tag{4}$$

Weaving the scaffold

When a descriptive hand motion ends, spatially proximate profiles are identified and joined together to construct a scaffold. In this process, longitudinal 3D curves are added to form an intercrossed net of 3D curves (Figure 3e). This net resembles the familiar hatching in traditional sketching techniques that designers use to quickly convey surfaces and curvature without elaborate shading. The density of the intercrossed curves may be appropriately adjusted so that the scaffold will be neither too dense nor too sparse.

Since the generated scaffold is implemented as a collection of generic 3D curve primitives, the user can operate on them as she would on pen-drawn 3D curves. By not construing the scaffold as a separate class of model, we avoid bloating our interaction vocabulary with primitive-specific interactions. Also, in our implementation, the scaffold maintains its original scale, so that the user can tap into her kinesthetic memory when working with these scaffolds. The scaffolds' visibility can be turned on and off as needed.

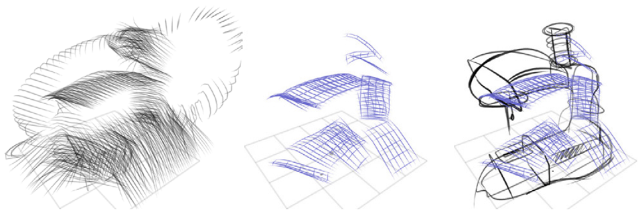

Figure 4. Our air scaffolding technique (which, in this case, was performed for 13 seconds) (a) extracts only the intended, descriptive shape information from the captured raw data and (b) visualizes it so that the user can (c) use it as a reference to author a high-fidelity 3D shape (a sewing machine).

Pen Sketching with Scaffolds

Having created air scaffolds, the user then defines sketch planes using the scaffolds as references and draws on them using pen-based 3D sketching (Figure 4c). To progressively express the overall shapes and finer features in 3D curves, the user iteratively adds and removes scaffolds and sketch planes as needed.

We note that our workflow shares the same underlying pen-based 3D sketching vocabulary as SketchingWithHands [16], in which the tick-based plane definitions enable quick and effective 3D sketching. Our workflow broadens the application domain of SketchingWithHands, which is limited to design conceptualization for handheld objects.

Creating sketch planes on scaffolds

The user performs one to three pen-tick gestures crossing existing 3D curves to define sketch planes (Figure 5) [16]. The intersections between the tick gestures and the 3D curves act as point constraints that the sketch plane must satisfy. Pen ticks can be performed on pen-drawn 3D curves and the scaffold's individual 3D curves alike. However, we assign higher priority to pen-drawn 3D curves so that when a tick gesture simultaneously crosses a pen-drawn 3D curve and a 3D curve on a scaffold, a point constraint will only be defined on the pen-drawn 3D curve. Intersection points between the sketch plane and any 3D curve that penetrates it are highlighted in yellow, enhancing the user's spatial awareness while the user changes views, draws 3D curves, and adjusts the sketch plane.

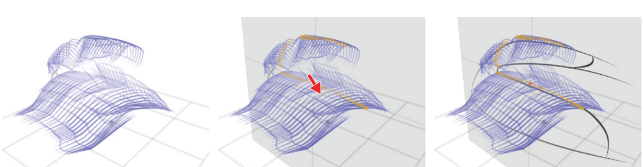

Figure 5. After (a) the user generates air scaffolds (blue) while 3D sketching an iron, (b) the user performs a tick gesture (red) on one of the curves of the air scaffolds to define a sketch plane. (c) The user sketches the side view of the iron while viewing the intersection points (yellow) between the sketch plane and the air scaffolds.

Erasing scaffolds

As the user progressively and iteratively adds more 3D curves, some scaffolds may lose relevance and simply get in the way. Thus, at any stage of our workflow, the user can conveniently remove parts of a scaffold or an entire scaffold (Figure 6). In implementing the scaffold eraser, we employed the intensifier concept proposed by Kim and Bae [16]. The default scaffold eraser removes an entire scaffold when the eraser comes into contact with any part of the scaffold, while the intensified eraser only removes the segments of scaffolds that the eraser directly touches.

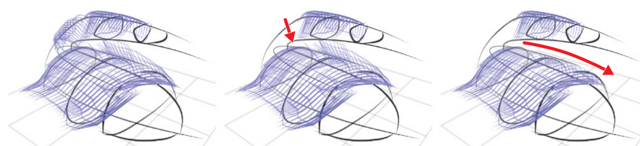

Figure 6. (a) In our workflow, the user can use the scaffold-only eraser to either (b) remove an entire group of air scaffolds or (c) trim only the unnecessary parts of the scaffolds along the path of the eraser.

Revisiting Workspace

At any time during our workflow, the user can generate additional scaffolds as needed. However, in doing so, the relative placement between the virtual 3D sketch and actual hand position can be lost due to the user's limited spatial memory [10]. To assist the user in finding the relative placement, we provide useful feedforward and feedback cues through responsive spangles [16].

Circular spangles with varying colors visualize the distance between the hand in the air and the 3D curves in the virtual space (Figure 7). As the hand approaches the 3D sketch and even moves through it, spangles first appear and then change colors to indicate varying degrees of proximity to the hand. Using these spangles, the user can quickly check the spatial relationship between the 3D sketch and the current hand position, revisit the previous workspace, and generate additional scaffolds.

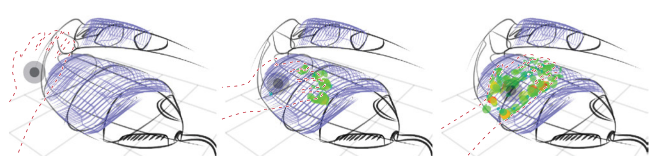

Figure 7. When the user's hand revisits the workplace, (a) concentric grey disks that indicate the center of the hand are displayed, and (b) when the hand approaches the 3D shape and makes contact, green responsive spangles are displayed at the contact points. Further penetration is visualized in yellow to red. (c) The user can utilize this sense of spatial relationship to iteratively add new scaffolds. (The user's hand—marked as a red dotted outline above—is not visualized in the system.)

IMPLEMENTATION

We implemented our workflow in a tabletop environment using a Wacom Cintiq 21UX graphics tablet and a Leap Motion hand-tracking sensor. Our configuration requires minimal hardware, which can be set up in any typical design workspace while also providing enough space between the tablet and the user for unconstrained hand motion (Figure 1).

Our software was written in Java using Leap Motion SDK, OpenGL, and the SketchingWithHands module [16]. All of the suggested techniques demonstrated and discussed hereafter were run in real time on a laptop with an Intel Core i7 2.70GHz CPU and Nvidia Quadro K2000M GPU.

DESIGN WORKFLOW SHOWCASE

We present a step-by-step design example to showcase our implemented workflow in a realistic design process. During the 3D sketch of a vacuum cleaner (Figure 8), which took 6 minutes and 38 seconds, 25 sketch planes were defined, taking an average of 5 seconds between drawing 3D curves, including camera navigation and sketch plane adjustments. The air scaffolding and pen-based 3D sketching were iteratively used to progressively flesh out the 3D shapes in accordance with the conventional product design process, where rough overall shapes are created first and details are added progressively.

We present a variety of sample 3D design concepts generated from our implemented workflow (Figure 9). Each 3D sketch took less than 10 minutes to complete and demonstrates the effectiveness of generating rough shapes using hand motion and sketching high-fidelity features using pen drawing.

The air scaffolding technique enables the creation of scaffolds that effectively describe the characteristic features of 3D objects. In drawing a relatively prismatic camera (Figure 9d), spreading the palm created flat scaffolds for the camera's casing, and extruding the curled palm created cylindrical scaffolds for the external microphone. In generating rough shapes for the skate (Figure 9b), a slightly closed palm expressed a gentle slope. When sketching the drone (Figure 9c), loosely closed hands were simultaneously extended in opposite diagonal directions to generate scaffolds for four thin arms of the drone. In sketching the more organic flower (Figure 9e), dynamically opening and closing the palm expressed curved surfaces with varying degrees of concavity. Such behaviors are in line with the observations of Holz and Wilson [10] and demonstrate that our air scaffolding worked in tandem with the designer's intuition.

While the scaffolds are not of high fidelity, the spatial information held by the scaffolds can effectively help pen-based 3D sketching produce more delicate 3D curves. In the case of the lamp object (Figure 9a), the scaffolds cover relatively small regions and do not conform tightly to the final outcome, but further 3D design ideas were developed through setting appropriate sketch planes using the scaffolds. In the case of the skate (Figure 9b), only the top and heel portions were expressed with the air scaffolds describing the basic scale and characteristic features, and other details, including the buckles, were drawn on top of the sketched 3D curves without the corresponding scaffolds. In the case of the flower (Figure 9e), only the pot, petals, and leaves were expressed using hand motion, while the stems were sketched later to join them together.

USER STUDY

We conducted a user study to see if our intended users could easily learn and use our implemented workflow and to evaluate the usefulness of air scaffolding in 3D sketching.

Participants

We invited 12 students (P1–P12) aged 20–24 (2 males and 10 females) who completed a basic drawing course at a department of industrial design in a university. All of them were educated in product design skills and were proficient in 2D digital sketching using a graphics tablet. All of the participants except P10 had experience using 3D modeling software, and none of them except P12 had experience in 3D sketching.

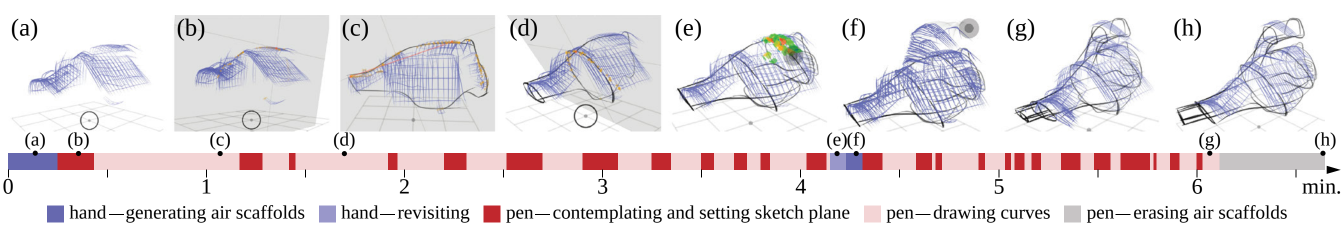

Figure 8. Outline of the entire working process for authoring a 3D sketch of a handheld vacuum cleaner: (a) creating rough overall shapes with air scaffolding, (b) seeking an appropriate viewpoint that captures the overall shapes and defining a sketch plane across the scaffolds, (c) drawing a dominant profile with pen sketching, (d) iteratively adding cross sections along the dominant profile by defining sketch planes and drawing curves on the planes, (e) revisiting the previous workspace (for about 5 seconds) based on the feedforward and feedback indications of the responsive spangles, (f) generating additional scaffolds needed to sketch the grip of the product, (g) drawing the dominant profile and section curves of the grip, and (h) trimming parts of the scaffolds that protrude from the finalized 3D sketch.

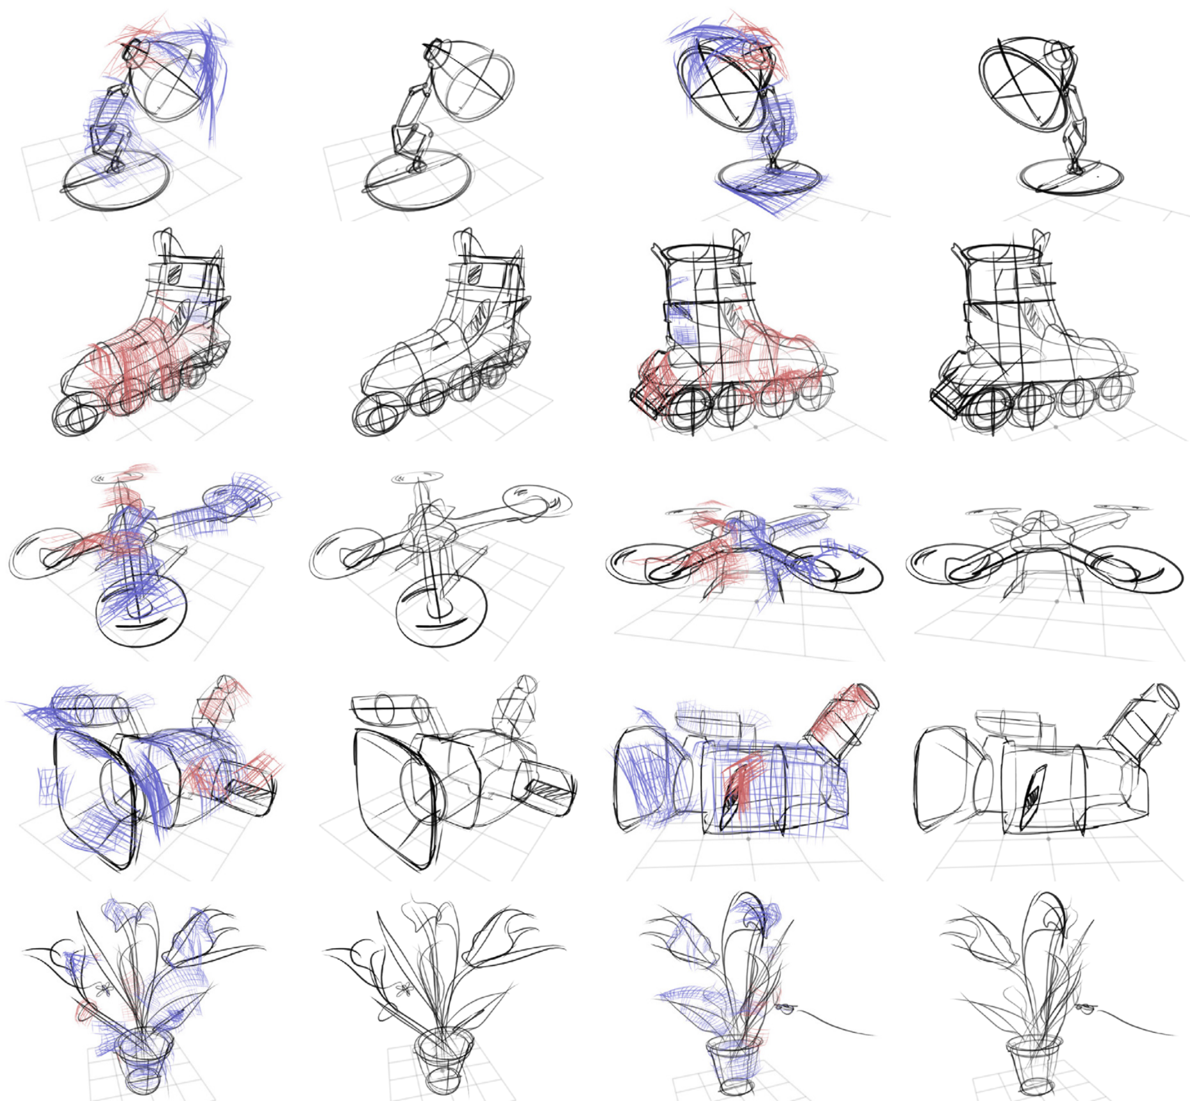

Figure 9. 3D sketch outcomes produced using our workflow and total time (h:mm:ss) taken. The red and blue scaffolds were created using the left and right hands, respectively. The scaffolds can be turned on and off.

Procedure

Each participant attended an orientation on our system on a one-on-one basis. One of us gave a 10-minute demonstration of the system and showed 3–4 sample sketches. We then held a 10-day sketching period, during which the participants could come in and use the system as many times as they wanted and freely use it to explore product ideas, regardless of time or the number of sketches. We provided a one-page cheat sheet with all of the functions and let the participants make comments or ask for help at any time. We collected the participants' comments, video recordings of the screen and the workspace, sketch outcomes, and the execution log of the authoring process throughout the study.

After the sketching period, we conducted in-depth interviews with the participants to discuss the pros and cons of the air scaffolding technique and suggestions for the system. To help the participants recall their experiences from the sketching period, we let them review their sketch outcomes and the video footage of their sketching process. In addition, we made them draw the same products twice with and without air scaffolding (in counterbalanced order) to help them identify the merit of air scaffolding in their workflow. At the end of the interviews, they completed a questionnaire on their preferences for each function of the system and the merits of air scaffolding using a 5-point Likert scale.

Results

All of the participants learned the functions of the system without difficulty and applied them in their design process. The participants visited our lab 3–5 times during the sketching period. The accumulated total usage time was 50 hours and 51 minutes, and 85 total concepts were produced (Figure 10). Since we allowed idea exploration without any time limit, some of the participants (P1, P3, P7, P8) became more skilled at using the system and immersed themselves in sketching for longer time spans. P1, who practiced voluntarily for the longest time, worked on a concept (P1-G) for over 2 hours during the last sketching session. We present noteworthy sketches from each participant in Figure 11.

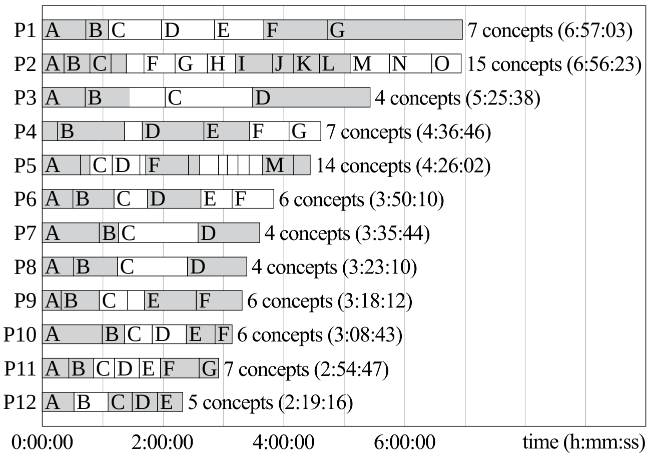

Figure 10. Usage time and number of generated concepts during the sketching period by participant. Visits are separated by shading and each concept is alphabetically coded.

The survey results (Figure 12) showed that the participants liked creating the scaffolds with hand motion (M = 4.00, SD = 0.60) and defining sketch planes on them (4.00, 1.35) in general. The most preferred function was toggling the scaffolds on and off (4.33, 0.78), which shows the natural workflow whereby participants received help from the scaffolds at the beginning and shed the aiding structure as sketching progressed. When and in what ways the scaffolds were preferable are further discussed in the next section.

The strongest merits of using air scaffolding were the time-efficiency (4.67, 0.65) and accuracy (4.42, 0.79) of defining scale and proportion. The ease (3.92, 0.29) and time-efficiency (3.83, 0.72) of creating a concept as well as the quality of the design outcomes (3.67, 0.78) were also evaluated highly. The merits of air scaffolding will also be further discussed together with the participants' qualitative feedback and sketched outcomes in the next section.

Figure 11. 3D design sketch outcomes produced by the participants using our workflow, and total time (h:mm:ss) taken.

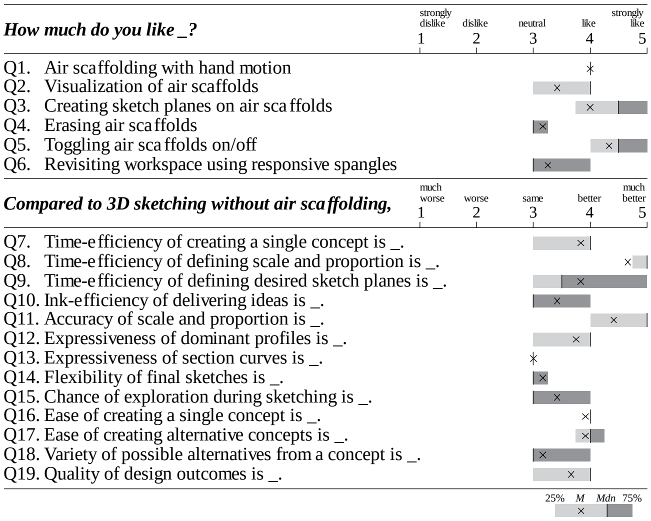

Figure 12. 5-point Likert scale evaluation by the participants.

ANALYSIS AND DISCUSSION

To better understand how air scaffolding facilitates the 3D sketching workflow, we first qualitatively analyzed the participants' feedback from the sketching sessions and interviews, and cross-checked the analysis with the sketched outcomes and survey results.

Air scaffolding was easy to learn.

As observed by Holz and Wilson [10], all participants were familiar with expressing shapes with hand motions and could immediately create air scaffolds: "*I don't think you need to learn to express shapes with hands*" (P5). The participants quickly found ways to incorporate air scaffolding into their unique workflow. In the first session, P8 had trouble using air scaffolds to draw a curved surface on the nose of a shoe (P8-A) but later authored various concepts (P8-B, P8-C, P8-D) that closely relied on the air scaffolds. P2 used air scaffolds to draw simpler geometry such as the projector (P2-B) at the beginning but later used them to construct more complex forms that needed to conform to the hand holding it, such as a hair dryer (P2-H) and a blender (P2-I).

Air scaffolding kick-started idea exploration.

The participants reported that air scaffolding was particularly useful during the initial stage when they had little or no idea as to what they wanted to draw: "*Like playing with a lump of clay, it gives me a lot of ideas about the 3D shape*" (P11). "*When I don't have a clear idea about the shape, it is difficult to draw a curve in 3D, but even then, I can still try this and that with scaffolding*" (P6). "*I can very quickly explore how to arrange components and in what proportions at the beginning of a design process*" (P11).

Air scaffolding enabled envisioning products in real scale.

The participants could easily express the approximate volume of the shape they desired in real scale using the sense of space: "*It is much more intuitive because I don't need to interpret the scale of the virtual space on the screen, and I can directly describe how big things are going to be*" (P10). Some of the participants created scaffolds in one-to-one scale to better capture the use scenario of the product: "*I feel like I can really grasp this camera* (P9-C)" (P9). "*The vase sketched without the scaffolds* (Figure 13a) *looks off when it is enlarged to real scale*" (P10). The survey also indicates that the scaffolds are useful for defining desired scale and proportion quickly (Q8, Q11), in support of these feedback.

Air scaffolds kept the 3D sketching in proper proportion.

The air scaffolds provided a consistent sense of 3D proportions that the participants could rely on from any viewpoint during 3D sketching: "*Without scaffolds, it is difficult to draw shapes that will look right from other directions*" (P4). "*When I use orthogonal cross-section sketches, I constantly have to update the proportions in other views when a sketch in one of them changes. But with air scaffolds, I can set and change key proportions right away*" (P11). "*I could not manage the overall proportion of the scooter without scaffolds* (Figure 13b)" (P11).

Air scaffolds eased locating sketch planes in any orientation.

The participants could find 3D constraints with which to define suitable 3D sketch planes from the approximate shape information that the air scaffolds contained: "*When I don't have the scaffolds, I have to define sketch planes in relation to the* x, y *and* z *axes only*" (P4). "*Scaffolds give me many reference points in the air that I can use*" (P12). P6 mostly used orthogonal sketch planes (P6-B) during the first session but later used a set of sketch planes with an organic flow (P6-F) based on the air scaffolds. Also, participants reported that the numerous intersection points between the scaffolds and a sketch plane (Figure 5b) were helpful in setting a new sketch plane: "*In an empty space, it's difficult to understand the position and orientation of a plane, but I can easily understand them when a plane has many intersection points with the scaffolds*" (P6).

Air scaffolds made 3D sketches minimal.

The air scaffolds helped the participants visualize what the final shape might look like from the beginning. As such, the participants could avoid unnecessary construction curves and focus on key characteristic curves, leading to more minimal sketches: "*It is easy to lose the sense of the whole when I'm drawing a part, but with the scaffolds, I can approximate the whole first and then start from the most important part*" (P12). "*Before, I drew curves in smaller segments because I couldn't have known what the entire shape would look like, but with scaffolds, I could draw curves more confidently knowing that they will be within the boundaries they need to be in*" (P11). "*I don't have to draw construction lines when I*

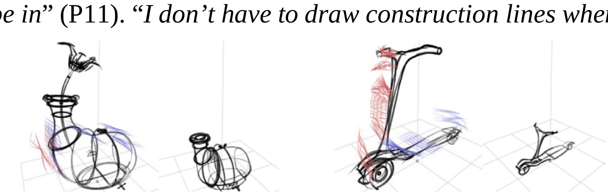

(a) vase (P10) (b) scooter (P11)

Figure 13. Products sketched by the participants with and without air scaffolding, showing that the scaffolds prevented unintended distortions in scales and proportions.

want to position a sketch plane in the middle of nowhere" (P3). These findings are consistent with the survey results, in that air scaffolding helped the participants express dominant profiles (Q12, M = 3.75, SD = 0.87). P10 drew noticeably fewer section curves as the sketching sessions progressed (P10-A~B, P10-D~E). Similar reductions of unnecessary curves were observed in others (Figure 14).

Air scaffolds had versatile utilities.

Participants also freely edited the air scaffolds to serve various needs. As the shape of the products matured, participants deleted parts or entire scaffolds that became unnecessary with the rough eraser: "*I only left behind parts of the scaffolds that could be used in drawing a new shape*" (P6). Many participants chose to hide the scaffolds after drawing the key curves to avoid visual clutter, while turning them back on temporarily as needed: "*When I forgot what the whole shape would look like, I turned it back on*" (P12). The higher preference for toggling the scaffolds on and off (Q5) over erasing them (Q4) supports this. Near the end of the drawing, some participants added finishing touches by trimming the edges of the scaffolds with the precise eraser, using the scaffolds like hatchings of the curved surfaces: "*If I leave behind the scaffolds, I can make it appear more volumetric without having to draw many lines*" (P8).

Air scaffolds stimulated new ideas.

The rough, approximate appearance of the air scaffolds could be interpreted as many shapes depending on the viewing direction, and this ambiguity spurred exploration and prevented fixation. The participants reported a phenomenon akin to *pareidolia*, wherein they saw different unexpected shapes in the air scaffolds, which helped them discover and develop new ideas: "*The air scaffolds are like the clouds. I sometimes found various shapes I didn't expect in the air scaffolds*" (P6). Some participants even intentionally utilized this phenomenon to search for various alternative design possibilities (Figure 15): "*I was pleasantly surprised that the same scaffolds could be used to draw so many different shapes. When I realized this, I tried to make out different shapes from the fuzzy scaffolds and develop them into completely different concepts*" (P2).

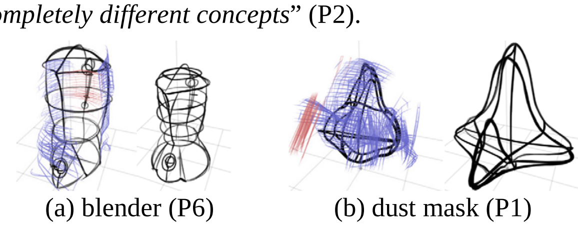

(a) blender (P6) (b) dust mask (P1)

Figure 14. Participants' sketches with and without air scaffolding—the scaffolds reduced unnecessary curves.

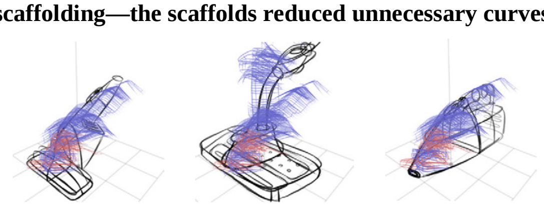

Figure 15. A participant (P2) decided to use the same scaffolds to draw different concepts of a vacuum cleaner.

LIMITATIONS AND FUTURE WORK

Because of the human body's limited mobility in different directions, the air scaffolds tended to be squashed in the anterior direction and lean toward the body. Such distortion is consistent with previous findings [10]. Despite this, the participants found the scaffolds suitable for their needs. They could simply compensate for the distortion by drawing properly proportioned sketches with the pen.

Some participants (P2, P4, P7, P10) pointed out that our prototype system only supports planar 3D curves and wished to draw free-space 3D curves with it. We conducted this study with the SketchingWithHands subsystem, but we foresee the air scaffolding technique being used with other 3D sketching systems. For instance, a head figure (P7-C) could be more effectively created using volume inflation techniques such as Teddy [11] and FiberMesh [18] alongside additional volumetric cues provided by air scaffolds, and non-planar curves could be drawn with ILoveSketch [2].

While the hand-tracking sensor we used had a range of about a meter, we may extend our technique to AR and VR, for which room-scale 3D authoring tools such as Tilt Brush [9] and Quill [19] are under active development. For example, a user can create large, rough volumes directly in space using air scaffolding while walking around in a virtual environment and then add fine details using perspective drawing with a pen and tablet. Such a system could enable the designing of larger products, as eagerly requested by some participants, without the problems such as fatigue and imprecision arising from lack of tactile feedback that perplex the current AR and VR tools [1].

CONCLUSION

Various computer graphics and 3D interaction techniques strive for more intuitive and expressive 3D modeling in design processes. In this study, we focused on the strengths and weaknesses of two salient input modalities and sought to develop a workflow in which the strengths of each are enforced and the weaknesses are complemented. Our contribution is in recognizing that 3D sketching has no technique analogous to 2D scaffolding in 2D sketching. 3D authoring utilizing hand motion was used to roughly and quickly generate the desired scale and volume, and pen-based 3D sketching was used to delicately specify finer features. From the extensive user study, we found that it is possible to design an interaction technique that coherently integrates different input modalities to enable rapid and high-fidelity 3D conceptualization in a progressive and iterative workflow that the designers could satisfactorily use in practice.

ACKNOWLEDGMENT

This research was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2013027260).

REFERENCES

1. 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* (CHI '17), 5643-5654.
2. 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* (UIST '08), 151-160.
3. 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* (UIST '09), 59-68.
4. Bruno R. De Araùjo, Géry Casiez, and Joaquim A. Jorge. 2012. Mockup builder: direct 3D modeling on and above the surface in a continuous interaction space. In *Proceedings of Graphics Interface 2012* (GI '12), 173-180.
5. Chris De Paoli and Karan Singh. 2015. SecondSkin: sketch-based construction of layered 3D models. *ACM Trans. Graph.* 34, 4, Article 126, 10 pages.
6. Chao Ding and Ligang Liu. 2016. A survey of sketch based modeling systems. *Frontiers of Computer Science* 10, 2: 985-999.
7. Julie Dorsey, Songhua Xu, Gabe Smedresman, Holly Rushmeier, and Leonard McMillan. 2007. The mental canvas: a tool for conceptual architectural design and analysis. In *15th Pacific Conference on Computer Graphics and Applications* (PG '07), 201-210.
8. 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, 1020-1032.
9. Google. 2016. Tilt Brush. https://www.tiltbrush.com
10. Christian Holz and Andrew Wilson. 2011. Data miming: inferring spatial object descriptions from human gesture. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '11), 811-820.
11. Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 1999. Teddy: a sketching interface for 3D freeform design. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (SIGGRAPH '99), 409-416.
12. Bret Jackson and Daniel 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.
13. Kiia Kallio. 2005. 3D6B editor: projective 3D sketching with line-based rendering. In *Proceedings of the 2nd Eurographics Workshop on Sketch-Based Interfaces and Modeling,* 73-79.
14. Olga A. Karpenko, John F. Hughes, and Ramesh Raskar. 2004. Epipolar methods for multi-view sketching. In *Proceedings of the First Eurographics Conference on Sketch-Based Interfaces and Modeling* (SBIM '04), 167-173.
15. Daniel F. Keefe, Robert C. Zeleznik, and David H. Laidlaw. 2007. Drawing on air: input techniques for controlled 3D line illustration. *IEEE Transactions on Visualization and Computer Graphics* 13, 5, 1067-1081.
16. 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* (UIST '16), 797-808.
17. Kevin T. McDonnell, Hong Qin, and Robert A. Wlodarczyk. 2001. Virtual clay: a real-time sculpting system with haptic toolkits. In *Proceedings of the 2001 symposium on Interactive 3D graphics* (I3D '01), 179-190.
18. Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. 2007. FiberMesh: designing freeform surfaces with 3D curves. *ACM Trans. Graph.* 26, 3, Article 41, 9 pages.
19. Oculus. 2016. Quill. https://www.oculus.com/story-studio/quill
20. Luke Olsen, Faramarz F. Samavati, Mario Costa Sousa, and Joaquim A. Jorge. 2009. Sketch-based modeling: a survey. *Computers & Graphics* 33, 1, 85-103.
21. Amit Pitaru. 2003. Rhonda. http://rhondaforever.com
22. 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* (CHI '01), 261-268.
23. Johannes Schmid, Martin Sebastian Senn, Markus Gross, and Robert W. Sumner. 2011. OverCoat: an implicit canvas for 3D painting. *ACM Trans. Graph.* 30, 4, Article 28, 10 pages.
24. Ryan Schmidt, Azam Khan, Gord Kurtenbach, and Karan Singh. 2009. On expert performance in 3D curve-drawing tasks. In *Proceedings of the 6th Eurographics Symposium on Sketch-Based Interfaces and Modeling* (SBIM '09), 133-140.

25. Ryan Schmidt, Azam Khan, Karan Singh, and Gord Kurtenbach. 2009. Analytic drawing of 3D scaffolds. *ACM Trans. Graph.* 28, 5, Article 149, 10 pages.
26. Jonathan Taylor, Lucas Bordeaux, Thomas Cashman, Bob Corish, Cem Keskin, Toby Sharp, Eduardo Soto, David Sweeney, Julien Valentin, Benjamin Luff, Arran Topalian, Erroll Wood, Sameh Khamis, Pushmeet Kohli, Shahram Izadi, Richard Banks, Andrew Fitzgibbon, and Jamie Shotton. 2016. Efficient and precise interactive hand tracking through joint, continuous optimization of pose and correspondences. *ACM Trans. Graph.* 35, 4, Article 143, 12 pages.
27. Edit Varga, Imre Horváth, Zoltán Rusák, Bram de Smit, and Han Broek. 2004. Survey and investigation of hand motion processing technologies for compliance with shape conceptualization. In *ASME 2004 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, 575-587.
28. Vinayak, Sundar Murugappan, HaiRong Liu, and Karthik Ramani. 2013. Shape-It-Up: hand gesture based creative expression of 3D shapes using intelligent generalized cylinders. *Computer-Aided Design 45, 2*, 277-287.
29. Gerold Wesche and Hans-Peter Seidel. 2001. FreeDrawer: a free-form sketching system on the responsive workbench. In *Proceedings of the ACM symposium on Virtual reality software and technology* (VRST '01), 167-174.
30. Ying Wu and Thomas S. Huang. 2001. Hand modeling, analysis and recognition. *IEEE Signal Processing Magazine* 18, 3, 51-60.
31. 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* (VRST '08), 223-226.