Window-Shaping: 3D Design Ideation by Creating on, Borrowing from, and Looking at the Physical World

Ke Huo, Vinayak, Karthik Ramani *

School of Mechanical Engineering, Purdue University
West Lafayette, IN 47907, USA
{khuo, vinayak, ramani}@purdue.edu

ABSTRACT

We present, *Window-Shaping*, a tangible mixed-reality (MR) interaction metaphor for design ideation that allows for the direct creation of 3D shapes on and around physical objects. Using the *sketch-and-inflate* scheme, our metaphor enables quick design of dimensionally consistent and visually coherent 3D models by borrowing visual and dimensional attributes from existing physical objects without the need for 3D reconstruction or fiducial markers. Through a preliminary evaluation of our prototype application we demonstrate the expressiveness provided by our design workflow, the effectiveness of our interaction scheme, and the potential of our metaphor.

ACM Classification Keywords

H.5.1 Information Interfaces and Presentation: Multimedia Information Systems; H.5.2 Information Interfaces and Presentation: User Interfaces; I.3.7 Computer Graphics: Three-Dimensional Graphics and Realism

Author Keywords

Design ideation; Sketch-based 3D modeling; Mixed reality; Augmented reality

INTRODUCTION

The physical environment often serves as a means for inspiring, contextualizing, and guiding the designer's thought process for expressing creative ideas. In early design processes, physical objects are frequently used as references to explore the space of novel designs [31, 35, 10]. Recent works [30, 36, 40] have shown that *through-the-screen* Augmented Reality (AR) and Mixed Reality (MR) can play a vital role in bridging the gap between the physical and digital worlds for creative expression of ideas. However, most of these approaches use the physical environment mainly as a dormant container of digital artifacts rather than as a source of inspiration for facilitating quick digital prototyping for design ideation. The key potential value that AR/MR systems brings to design, is the integration

*School of Electrical and Computer Engineering (by courtesy)

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

TEI '17, March 20-23, 2017, Yokohama, Japan

© 2017 ACM. ISBN 978-1-4503-4676-4/17/03...\$15.00

DOI: <http://dx.doi.org/10.1145/3024969.3024995>

Figure 1. Users simply draw a curve on the screen (a), that is mapped to a 3D planar curve using the point cloud (b). The 3D curve is inflated into a 3D model (c). Users manipulate the shapes through a multi-touch interaction scheme (d). *Window-Shaping* enables quick creation of virtual artifacts for augmenting the physical environment by borrowing dimensional and textural attributes from objects (e).

of *reflection-in-action* [31] (creating on the physical world), *design inspiration* [11] (borrowing from the physical world) and *reflection-on action* [12] (looking at the physical world). In this paper, we explore this value through re-purposing the physical environment as a reference, context, and source of inspiration for quick idea generation in early design.

We present *Window-Shaping*, an approach that integrates sketch- and image-based [21, 24] 3D modeling approaches within a mixed-reality interface to develop a new design workflow (Figure 1). Using the Google Tango device, *Window-Shaping* leverages the RGB-XYZ (i.e. image and point-cloud)

representation of a scene allowing users to create planar curves on physical surfaces and *inflate* them into 3D shapes. Using *Window-Shaping* we demonstrate design scenarios including the use of everyday objects and low-fidelity mock-ups as design references, and exploration of novel designs by combining physical references from multiple sources. *Window-Shaping* both complements and extends existing approaches [36, 18, 38] through presenting a new interaction metaphor wherein a new virtual 3D object is created as an extension of its physical context without the need for reconstructing the 3D model of the physical scene. We make two contributions:

Tangible In-Situ Design Workflow: We offer a novel combination of an existing modeling scheme with synchronized RGB-XYZ information to enable creative design exploration with the physical environment in context.

Dimensionally consistent and visually coherent design: We offer the capability of creating and visualizing 3D shapes directly on the surface of any object with the desired dimensions and locations. Further, by mapping the background texture of the users' sketch inputs, we allow users to re-purpose existing textures in new creations.

RELATED WORKS

Our main focus in this work is on a novice-friendly interaction metaphor allowing design directly in 3D for *early stage design ideation*. We discuss the relevant works selectively.

Design in Context

Most of the existing approaches for digital design focus on modeling from scratch rather than on allowing novice users to draw inspiration from existing physical artifacts - a frequent activity in creative idea generation. Importantly, we use the physical environment as a dimensional and visual reference. Although many works infer the design context in various forms, such as static images [19, 39, 4], registered bundle photographs [25], and 3D models [33, 36, 5, 8], they do not allow for modeling directly in 3D space. *Window-Shaping* focuses on creation of new geometry in an augmented 3D scene through a hand-held interface. Our approach uses RGB-XYZ representation of the physical environment.

Content Creation with AR

Virtual content creation using an AR-based system has also been explored. Previous works integrated instrumented tangible tools for operations such as modifying virtual models [1]. Nuernberger et al. [22] interpreted 2D drawing annotations using cues from 2D images and 3D geometry. *SnapToReality* [23] took a step further by extracting 3D edge and planar surface constraints from the environment and using them for precise alignment. Lau et al. [18] attached fiducial markers onto physical primitives and stamped the corresponding virtual shapes together to create shapes. Leveraging the merging mobile AR techniques, Xin et al. [38] presented *Napkin Sketch* as a system for creating 3D wire-sculptures on a napkin. Recently, *MixFab* [36] integrated scanned 3D model towards design using mid-air gestures in MR environments. While we draw inspirations from these approaches, our aim is to allow users to design directly on physical objects without being constrained by set-ups or being limited by the mobility.

Geometric Design with Hand-Held Devices

Previous explorations on 3D sketching using hand-held devices tended to separate the input devices from the display [28, 7]. Following a similar approach, Vinayak et al. [34] and Piya et al. [26] demonstrated systems that utilized smartphones as multi-touch controllers. Further, past works approached an immersive design environment by incorporating projection-based [15] or see-through head mounted displays [29]. Other works such as *(T)ether* [17] leveraged both touch inputs and mid-air gestures for modeling using a tracking infrastructure in conjunction with instrumented wearables and a tablet. *Paper3D* drew inspiration from papercraft and used multi-touch gestures for casual 3D modeling. In these works, the 3D shapes were created in either an empty physical space or in a virtual environment thus neither the dimensionality nor the visual appearance might pertain to their designated environment. While *Napkin Sketch* and *Second Surface* [14] started to merge the mobile AR technique, the physical reference in their works was limited to smaller working volumes and their outcomes were 3D wire or drawings. We concentrate on enabling the creation and editing of 3D models on arbitrary physical surfaces using a simple multi-touch interaction scheme.

WINDOW-SHAPING

The *Window-Shaping* interface comprises a hand-held Google Tango device, that serves as a local interface between the physical environment and the user. The simultaneous localization and mapping (SLAM) algorithm available with the Tango API allows for the acquisition of a point-cloud of the scene with respect to the global coordinate system. The resulting RGB-XYZ data allows users to implicitly define planes on any physical surface by simply drawing on the physical scene. Any touch input on the device screen can be *unprojected* on the physical environment to obtain a 3D point along with its normal in the global (world) coordinate system (Figure 1). This essentially helps users define a plane at any recognized point on a physical surface. Below, we describe the design goals, modeling metaphor, and user interactions.

Design Choices

We build towards the broad theme of "*effortless integration of physical objects into the design process*" [37]. The primary objective of our work is to support quick design ideation by allowing users to (a) quickly create 3D geometry in reference to physical artifacts, (b) borrow shape and appearance from physical artifacts to re-purpose them for design exploration, and (c) inspect the virtual artifacts in the physical context from different views in order to make design modifications. Our interface design choices are described below.

Appropriate Use of Interaction Modality:

The use of Google Tango tablet allows for both spatial (3D) and multi-touch (2D) interactions. While touch-based interactions allow for precise control and 2D sketching operations, spatial mobility allows for *reflection-on-action* by enabling users to inspect their creations from multiple views with respect to the physical environment. Thus, we use well-established multi-touch interactions for enabling content creation, editing, and rigid transformations. For inspection, we

Figure 2. Geometric primitives with different inflation functions in *Window-Shaping* include: a. Circular, b. Conical, c. Tapered, and d. Linear.

make use of the natural spatial movement. With the augmentation from the see-through video and captured RGB-XYZ information, the traditional 2D interactions go beyond planes, and are enhanced with a third dimension. Further, this three-dimensional extension of touch interactions provides a tangible and immersive experience for the design ideation.

Consistent dimension and appearance:

The appearance of designers' creations is an important factor in reflecting their intent. Our metaphor enables users to both use physical objects as contextual references as well as re-purpose them at different physical locations. Further, our approach allows copying the texture of physical references for a consistent rendering of the newly created 3D shapes. Further, the editing operations are designed such that the changes maintain the consistency of visual appearance.

Geometric Modeling Scheme:

We aim for flexibility in terms of the expressive power of the modeling scheme while retaining the simple interactions for shape creation. Multi-touch inputs naturally allow for 2D curve input. Thus, we employ a *sketch-and-inflate* modeling scheme in *Window-Shaping*. First, a user sketches the silhouette (and holes) of the shape with which a closed mesh is generated. Then we inflate the mesh using a distance transform function [24] to obtain an inflated 3D shape. We provide users with four inflation functions (primitives) for expressive shape creation (Figure 2). The main advantages of this approach are that it: (a) has been demonstrated to be particularly simple for novice users [13], (b) allows for creation of complex topological structures with a simple set of interactions, and (c) has a simple and natural 2D parametrization that allows for texture mapping.

User Interactions

Shape Creation and Editing

Projective Sketching: *Window-Shaping* allows for direct one finger drawing on the tablet screen. Once finalized, the sketched curve is mapped on the physical scene and is converted and rendered in the scene as a 3D inflated mesh (Figure 3(a, b)). The first curve drawn by the user is by default the boundary curve. Multiple *hole* curves can then be drawn inside the boundary.

Placing Curve Templates: As an alternative to direct drawing, we also provide a set of curve templates (Figure 3(c,d)). Users can simply place the selected curve on any surface of a physical or virtual object using a single-tap gesture. The curve is placed on a fitted 3D plane around the single-tapped location. The curve template feature allows for quick exploration of complex ideas with minimal interaction effort.

Capturing Outlines: *Window-Shaping* also allows users to extract the outline of the object from the scene in the *image space*.

Figure 3. The user sketches boundary and hole curves on a physical object, edits the sketched curves to add local details(b), and obtains an inflated circular shape (b). The user creates a tapered inflated shape with a template, inflates it, patterns it (c), explores complex features(d).

Users draw a region of interest (ROI) which is automatically converted into a contour using the GrabCut algorithm [27]. This enables users to directly employ the visual representation (outline shape and texture) of a physical object and re-purpose it in 3D form in their own designs.

Editing Curves: Using the over-sketching interaction technique [3], we provide simple and quick curve editing (Figure 3(a)), which enables users to add details and improve the appearance. Moving the tablet closer to a desired region allows for precise virtual operations in screen-space. On the other hand, moving away from a physical surface allows for a better overview, which is valuable for coarse operations such as placing shapes and curve templates on desired locations.

Inflating and Deflating: We implement a three-finger gesture for inflating or deflating a 3D mesh. Here, pinching (bringing fingers closer) effects in pulling the shape out of the screen as well as spreading (moving fingers apart) result in pushing the shape into the screen.

Manipulating Shapes

Rotations & Scaling: Two-finger rotate and pinch/spread are used for rotating and scaling the shape respectively. These gestures can be applied either directly to the 3D shape or to the underlying curve of the shape. The two-finger interaction constrains all rigid transformation to the plane of the curve.

Translation: The in-plane translation is performed by dragging a shape using one finger. This allows for precise placement of the shape on the plane defined by its underlying curve. In order to provide consistent dimensional perception, we project the finger movement onto the underlying plane instead of using constant mapping between the pixel space and the physical space.

Placement: Shape placement allows users to directly transfer a selected 3D shape to any point in the scene by using a one-finger tap gesture. Here the 3D shape is both translated to the specified point and re-oriented along the normal at this point. Here users can place a new virtual object on the physical scene as well as on an existing virtual object. This maintains a perceptual depth consistency during interactions.

Auxiliary Operations: In addition to geometric operations, we provide operations such as copying/patterning and deleting a shape. Users can select and make a copy/pattern of the shape by using a single-tap gesture.

Figure 4. Illustration of plane inference: Un-projection of a 2D drawing results in a discontinuous curve (red) and a projection on the inferred plane(green)

Appearance Control

During the over-sketching operation, we automatically update the texture image to maintain the visual consistency. We also provide the option to explicitly update the texture during rigid transformations. This is helpful to users when they are experimenting with different backgrounds for the same shape.

IMPLEMENTATION

Hardware & Software

Our hardware comprises a Google Tango 7 inch tablet with the NVIDIA Tegra K1 processor and 4GB RAM, running Android 4.4 KitKat OS. The tablet captures the RGB image (60Hz) and depth data (5Hz) from the built in 4MP color camera and depth sensor respectively. The Tango SDK [9] provides functionality for synchronizing these two cameras, allowing us to compute a point cloud (XYZ) of the scene such that each point is mapped to a unique pixel in the RGB image. We prototyped our metaphor using the Android SDK and the geometric modeling methods in C++ using the Android NDK (JNI) with OpenGL Shading Language for rendering.

2D Curve Processing

We require all curves to be closed, oriented, and preferably smooth while preserving the features. To meet these requirements, we first apply an exponential smoothing filter [32] to each point on the curve as the user is drawing them. We then check the curve for closure based on the distance between the end-points, discarding an open curve as an invalid input. For a closed curve, we perform an equidistant curve re-sampling [16] and orient the boundary curve counter-clockwise (i.e. positive area) with the holes oriented clockwise (negative area).

3D Planar Curve Computation

Given the processed curve on the screen, we first query the 3D points corresponding to each curve-point. For each 3D point we compute its normal by fitting a plane using its neighborhood. Based on the standard deviation of the distances between adjacent points on the 3D curve, we categorize the curve as either continuous or discontinuous. For a continuous curve, we estimate its plane by averaging the position and normals of these points. This, however, results in unpredictable planes for discontinuous curves (Figure 4). In this case, we first divide the curve into segments belonging to the same plane using euclidean distances and normal differences and then select the largest segment to identify the plane.

Figure 5. Creating a texture: (a) projecting the bounding rectangle (blue) of the 3D planar curve (red), (b) image skewing, (c) rotation correction, and (d) image cropping.

Figure 6. Furniture design: A virtual side-table is created by borrowing the texture from a physical table (a, b). The surrounding objects are then used to explore the lamp design (c, d) and *GrabCut* is applied to capture the outline and texture(e, f) to form a decorative object(g, h)

Mesh Generation

Given the processed boundary and hole curves, the mesh generation is performed in three steps: (a) computing two symmetrically aligned open (*half*) meshes bounded by curves (boundary and holes) through constrained delaunay triangulation (CDT), (b) topologically stitching these two open meshes to create a closed mesh, and (c) inflating the *top half mesh* using the distance transform function [24]. We implement CDT using the *poly2tri* library [6]. For the round, conical and tapered primitives, we sample the interior region of the curve with a uniform equilateral point configuration and add the sample as Steiner points to obtain a regularly sampled triangulation. Further, this modeling scheme has a simple and natural 2D parametrization which allows for texture mapping.

Texture Computation

We implement texture generation using *openCV* in four steps (Figure 5). We first compute the bounding rectangle of the 3D planar curve and project the bounding box on the image space. Then, we apply a skew transformation on this image with the constraint that the projected bounding rectangle is axis-aligned with the image. We rotate the skewed image to correct the residual angle between the rectangle and the image. Finally, we crop the image using this projected bounding rectangle to obtain the texture image.

USE CASES

The design work flow and interactions in Window-Shaping, can potentially be adapted to different kinds of design contexts. Below, we identify four such design patterns.

Figure 7. Chair armrest design: A truss-like shape is created on a metal shelf and placed as an armrest on the sides of a chair (a). A box is used here to appropriately position and orient the armrest with respect to the seat (b). Using a template (c), the back-rest is re-designed (d)

Figure 8. Creature design: The eyes (a), limbs (b), and body (c) are created using a helmet, a trash can, and a piece of white paper as visual references respectively. They are then assembled on a table (d) and the details (the scales of the creature) are created from a mat sheet and patterned on the body (e, f).

Designing on Physical Objects: The most important design capability offered by *Window-Shaping* is creating new geometric features on existing objects. These existing objects can be both physical and virtual objects. For instance, in an interior design scenario, a user could add complementary features to a piece of furniture (Figure 1, Figure 6(a,b)) and also create virtual additions to the scene by adding new assemblies to the surrounding area (Figure 6(g, h)).

Re-purposing Physical Objects: By re-purposing, we mean the use of both the shape and the appearance of a physical object to create a new design feature. The aforementioned GrabCut algorithm allows users to capture outlines of existing objects from the scene for shape creation. The captured outline shape and texture serve as design inspiration for direct use in an existing mixed-reality scene (Figure 6(e, f)).

Physical Objects as Spatial References: In situations where users desire to *fill in a blank space* to augment a physical product, it can be helpful to use a physical object to define a reference plane (Figure 7). Using objects as references enables a tangible and spatially coherent way of designing in context.

Physical Objects as Visual References: The appearance can serve both aesthetic and functional purposes (such as material specification). In *Window-Shaping*, users can experiment with the appearance of a 3D model. Such experiments can be performed either by transferring the virtual shape to a new location and re-texturing or by simply changing the background texture of a sketched curve (Figure 8).

USER TRIALS

We conducted a preliminary evaluation for eliciting user feedback regarding the utility, experience, potential, and limitations of our workflow in creative design activity.

Participants

We recruited 8 (2 female, 6 male) participants (26 – 30 years old) with 5 engineering students and 3 students from non-engineering fields (science, management, etc). Of these, 4 participants had no prior knowledge of AR or MR interfaces, 3 were familiar with the concept of AR/MR, and 1 had used AR interfaces for gaming. Three participants had no prior knowledge of computer-aided design (CAD) or 3D modeling.

Procedure

We conducted user trials (45 – 60 min. per trial) with the *Window-Shaping* prototype application. We first introduced participants to the broader idea behind *Window-Shaping* and demonstrated the user interactions in *Window-Shaping* through practical use-cases (10 min.). Following this, we invited the users to perform four design tasks. We gave the users the option to either use a capacitive stylus or direct finger touch for tablet interactions.

(T1) Designing with physical mock-ups where the participants were given a cuboidal box and were asked to create the face of a creature. **(T2) Re-purposing objects** through which we introduced the grab-cut feature to the participants and asked them to create a 3D new part using *GrabCut* and add it to the face. **(T3) Using an object as a spatial reference** the participants were asked to design the handle of a chair using a cuboidal for placing the handle on the sides of the chair. **(T4) Using objects as visual references**, the users added details to the back rest and seat and explored the texture of the handles.

The participants were allowed to move around in the surrounding environment, create their feature on an arbitrary object and then transfer the feature back on the design. While we guided the participants through the interactions, we encouraged them to define their own strategy for completing the design tasks. At the end of the tasks, they were asked, through answering a questionnaire, about their experience in using the interface (Figure 9(a, b, c)). We also asked them to explain their reasons along with the Likert scales. Moreover, we asked three open-ended questions regarding the usability, potential use scenarios, and desirable capabilities.

Findings

Although we constrained the design tasks, we found that the resulting creations (Figure 9(d)) had reasonable diversity across users. Most of the users were able to quickly understand the modeling mechanism and successfully perform the trial tasks within the given time. Below, we discuss the main insights we gained from our observation and the feedback from the users.

Validation of Interactions

All participants responded favorably to the ease of use offered by direct drawing and creating fine details and a majority preferred drawing more than using available curve templates (Figure 9(a)). Interestingly, there was a distinction between the engineering and non-engineering students in terms of the

Figure 9. User feedback on interactions (a, b) and overall experience (c), and designs generated in trial tasks (d).

curve creation. Engineering students expressed a need for interactions such as mirroring, symmetric curve creation and editing. A user suggested: “to enter some numeric values of the extrusion lengths”. In contrast, one non-engineering user commented on the curve templates provided: “I did not use these features. It was intuitive to create something by myself”.

Most users found the interactions for placing (one finger tap) and manipulating (two finger rotate-scale) the curves and shapes to be intuitive. A user noted: “Very similar to existing gestures on smartphones and tablets, so very intuitive”. Regarding shape inflation, while 6 out of 8 users agreed on the intuitiveness of the 3-finger gesture, 2 users expressed negative response to it. According to one user, “It was intuitive, but the direction could be inverted”. This was contrary to our assumption in which pinching increased the inflation and spreading decreased the inflation magnitude similar to in the physical action of deforming clay.

Mobility & Fatigue

One of the most important aspects enjoyed by all users was the capability to move around in physical space during the design process. This mobility offered the flexibility for users to borrow geometric and visual features from various locations (user comment: “this help(s) you also to create shapes everywhere at any time”). During the trial tasks, we observed that the users walked frequently within our constrained space to look for existing objects as references, to perform spatial inspection from different angles, and to design new shapes on reference planes. Further, some of the users naturally moved the device close to the object while creating detailed features on the curves. One user said: “Bringing it [the tablet] closer did help a lot.” This strongly suggests a positive outcome of our approach towards enhancing design ideation by increasing the reflection in and on the action.

While we expected this outcome, we were also interested in knowing how this mobility in our interface affected user fa-

tigue. Six users agreed that using our system did not cause physical stress or fatigue that could adversely affect the design task itself. We believe this to be a result of the balance between the size ($119.77 \times 196.33 \times 15.36 \text{ mm}$) and the weight (0.82 lbs) of the Tango tablet [9]. However, we did notice that our system caused fatigue specifically in the process of adding details to the curves for some users. A female user mentioned: “While getting precise details for a long time it may cause the arm to be a bit tired.”. Last, although we offered the users two choices between using a stylus and their fingers for touch inputs, we got no particular comments on this through the survey.

Design Workflow

Users responded favorably in terms of the expressiveness, engagement, and enjoyment provided by *Window-Shaping*. In particular, the user feedback strongly validated our primary goal – quick design ideation through dimensionally consistent and visually coherent shape modeling. In the given design contexts, all users except one agreed that our geometric primitives were expressive and allowed for sufficient scope for creative exploration. As expected, users generally favored the fact that our system allows capturing the background textures and can be used in real time. A user reported “[without using our system] it will be difficult specially for getting textures or for placing the features on the desired position.” Also, users tended to inspect the dimensionality of the creations from different perspectives. One user stated: “The scale of the object with respect to the surroundings could be easily found.”

Utility & Potential

Users confirmed *Window-Shaping*’s utility as a quick ideation tool both for individuals and design teams. As one user said: “The tool is definitely very useful for tinkering and showing new ideas to others.” In particular, users with prior CAD experience appreciated the value of designing with a real scene in context. One user commented: “You can visualize the components in the real world much more clearly than any CAD

software.” Another user pointed to the potential for collaborative design: “*I can see this being very useful in design collaboration where several people work on it at the same time. Very useful in modifying existing products and brainstorming what features to add.*”. Participants also proposed some novel applications such as “*It’s useful for topography when we need to study the landforms of some region*” and “*design your garden before planting*”.

Discussions & Limitations

User Interaction: Drawing while holding the device resulted in a lack of precise control. The participants found the curve editing to be more challenging in comparison with tasks such as placement. This primarily affected the addition of curve details. We believe this can be addressed through filtering out hand jitters by using the device IMU data and the camera image. Long term usage of the device may cause potential fatigue. Nonetheless, depth sensing technologies are emerging with more light-weight mobiles, such as *Lenovo Phab 2 Pro* [20]. Further, the texture capturing could be better implemented by allowing users to control the object segmentation. While users liked the idea of using physical objects as references, they pointed out a need for a mechanism to explicitly define a reference plane. This issue can be easily addressed by introducing *virtual plane* and *3D widget* elements into our interface. While previous work has shown better drawing accuracy for styli in comparison to fingers [2], it will be worthwhile to conduct more controlled experiments with a mobile system such as ours in terms of accuracy, fatigue, and stimulation.

Advanced AR Environment: While the robustness of the tracking provided by the Tango was impressive, our current implementation did not allow for object awareness. This resulted in shifting of the users’ creations due to the frequent device movement. Introducing object tracking will reduce the tracking artifact when dealing with featureless scenes. Even though our interface offers mobility in a large indoor environment, it is currently difficult to create or borrow geometric and textural features from smaller objects. At this stage, the point cloud quality is limited by the low accuracy of the depth camera. Another issue in our current implementation is occlusion management with real objects. To address this issue, usually a fine model of the object and a high resolution of the depth data are required. With constant hardware improvements to the depth camera and mobile computational power in the future, we believe this issue will be easier to resolve.

Modeling Scheme: We restricted our implementation to the use of planar curves for inflating shapes. Although the current modeling scheme allows for reasonable expressive capabilities, there is a need to investigate the usage of 3D curves as well as different shape representations such as swept primitives, skeleton-based models, and volumetric representations. Towards a more complete and refined application, additional standard operations such as undo and redo will certainly enhance the modeling work-flow. Further, improvements in point-cloud acquisition will also allow us to extract and use manifold constraints (e.g. corners, edges, curvature) from the environment. The assembly mating relationships and kinematic constraints could then be leveraged for designing functional objects.

CONCLUSIONS

Window-Shaping reveals an untapped design space that emerges from the combination of multi-touch interactions, sketch-based geometric design, and mixed-reality interface towards bridging the gap between physical and digital space in early phase design exploration. Given the users’ positive reactions, we believe that the proposed concept has potential towards a richer space of MR-based design work-flows for advanced in-situ modeling, collaborative idea generation, and fabrication-aware design. Although through the preliminary evaluations, we obtain overall positive feedback on the interface, the limitations of the current implementation need to be addressed before *Window-Shaping* reaches its broader potential. Based on the findings from our preliminary evaluations, we plan to improve the interaction metaphors against the jitter, accuracy and fatigue issues, and adding advanced geometric features into the existing modeling approach. Further, we plan to study how experience, performance, and creative outcomes will change with respect to different user groups such as artists, engineering designers, and young participants. Finally, it will be worthwhile to find how the interactions behind *Window-Shaping* could be extended to applications in domains such as architecture, education, animation, and engineering design and analysis.

ACKNOWLEDGMENTS

We thank the reviewers for their valuable feedback and comments. This work was supported partially by the NSF Awards No. 1538868 (CMMI) and 1632154 (PFI). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

1. Ryan Arisandi, Mai Otsuki, Asako Kimura, Fumihisa Shibata, and Hideyuki Tamura. 2014. Virtual Handcrafting: Building virtual wood models using ToolDevice. *Proc. IEEE* 102, 2 (2014), 185–195.
2. Sriram Karthik Badam, Senthil Chandrasegaran, Niklas Elmquist, and Karthik Ramani. 2014. Tracing and sketching performance using blunt-tipped styli on direct-touch tablets. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces*. ACM, 193–200.
3. Thomas Baudel. 1994. A mark-based interaction paradigm for free-hand drawing. In *Proc. of UIST*. ACM, 185–192.
4. Tao Chen, Zhe Zhu, Ariel Shamir, Shi-Min Hu, and Daniel Cohen-Or. 2013. 3Sweep: Extracting Editable Objects from a Single Photo. *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH Asia* 2013 32, 6 (Nov. 2013), 195:1–195:10.
5. Xiaowu Chen, Bin Zhou, Feixiang Lu, Lin Wang, Lang Bi, and Ping Tan. 2015. Garment modeling with a depth camera. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 203.

6. Vid Domiter and Borut Žalik. 2008. Sweep-line algorithm for constrained Delaunay triangulation. *International Journal of Geographical Information Science* 22, 4 (2008), 449–462.
7. Michele Fiorentino, Raffaele de Amicis, Giuseppe Monno, and Andre Stork. 2002. Spacedesign: A mixed reality workspace for aesthetic industrial design. In *Proceedings of the 1st International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, 86.
8. Madeline Gannon, Tovi Grossman, and George Fitzmaurice. 2015. Tactum: A Skin-Centric Approach to Digital Design and Fabrication. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1779–1788.
9. Google. 2016. Tango Developer Overview. (Jan 2016). <https://developers.google.com/tango/developer-overview>
10. Nicolai Brodersen Hansen and Peter Dalsgaard. 2012. The productive role of material design artefacts in participatory design events. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*. ACM, 665–674.
11. Scarlett R. Herring, Chia-Chen Chang, Jesse Krantzler, and Brian P. Bailey. 2009. Getting Inspired!: Understanding How and Why Examples Are Used in Creative Design Practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 87–96.
12. Caroline Hummels and Joep Frens. 2009. The reflective transformative design process. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2655–2658.
13. Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 2007. Teddy: a sketching interface for 3D freeform design. In *Acm siggraph 2007 courses*. ACM, 21.
14. Shunichi Kasahara, Valentin Heun, Austin S Lee, and Hiroshi Ishii. 2012. Second surface: multi-user spatial collaboration system based on augmented reality. In *SIGGRAPH Asia 2012 Emerging Technologies*. ACM, 20.
15. Daniel F Keefe, Daniel Acevedo Feliz, Tomer Moscovich, David H Laidlaw, and Joseph J LaViola Jr. 2001. CavePainting: a fully immersive 3D artistic medium and interactive experience. In *Proceedings of the 2001 symposium on Interactive 3D graphics*. ACM, 85–93.
16. Per-Ola Kristensson and Shumin Zhai. 2004. SHARK²: A Large Vocabulary Shorthand Writing System for Pen-based Computers. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology (UIST '04)*. ACM, New York, NY, USA, 43–52.
17. David Lakatos, Matthew Blackshaw, Alex Olwal, Zachary Barryte, Ken Perlin, and Hiroshi Ishii. 2014. T(Ether): Spatially-aware Handhelds, Gestures and Proprioception for Multi-user 3D Modeling and Animation. In *Proceedings of the 2nd ACM Symposium on Spatial User Interaction (SUI '14)*. ACM, New York, NY, USA, 90–93.
18. Manfred Lau, Masaki Hirose, Akira Ohgawara, Jun Mitani, and Takeo Igarashi. 2012. Situated Modeling: A Shape-stamping Interface with Tangible Primitives. In *ACM Conference on Tangible, Embedded and Embodied Interaction (TEI '12)*. 275–282.
19. Manfred Lau, Greg Saul, Jun Mitani, and Takeo Igarashi. 2010. Modeling-in-context: User Design of Complementary Objects with a Single Photo. In *Proceedings of the Seventh Sketch-Based Interfaces and Modeling Symposium (SBIM '10)*. Eurographics Association, 17–24.
20. Lenovo. 2016. Phab 2 Pro. (Jan 2016). <http://shop.lenovo.com/us/en/tango/>
21. Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. 2007. FiberMesh: Designing Freeform Surfaces with 3D Curves. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07)*. ACM, New York, NY, USA, Article 41.
22. Benjamin Nuernberger, Kuo-Chin Lien, Matthew Turk, and others. 2016a. Interpreting 2D gesture annotations in 3D augmented reality. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 149–158.
23. Benjamin Nuernberger, Eyal Ofek, Hrvoje Benko, and Andrew D Wilson. 2016b. SnapToReality: Aligning Augmented Reality to the Real World. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1233–1244.
24. Luke Olsen, Faramarz Samavati, and Joaquim Jorge. 2011. NaturaSketch: Modeling from Images and Natural Sketches. *IEEE Comput. Graph. Appl.* 31, 6 (Nov. 2011), 24–34.
25. Patrick Paczkowski, Min H Kim, Yann Morvan, Julie Dorsey, Holly E Rushmeier, and Carol O'Sullivan. 2011. Insitu: sketching architectural designs in context. *ACM Trans. Graph.* 30, 6 (2011), 182.
26. Cecil Piya, Vinayak, Yunbo Zhang, and Karthik Ramani. 2016. RealFusion: An Interactive Workflow for Repurposing Real-World Objects towards Early-stage Creative Ideation. In *Proceedings of Graphics Interface 2016*. ACM.
27. Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. 2004. "GrabCut": Interactive Foreground Extraction Using Iterated Graph Cuts. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 309–314.
28. Emanuel Sachs, Andrew Roberts, and David Stoops. 1991. 3-Draw: A Tool for Designing 3D Shapes. *IEEE Computer Graphics and Applications* 11, 6 (1991), 18–26.

29. 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*. ACM, 261–268.
30. Rajinder S. Sodhi, Brett R. Jones, David Forsyth, Brian P. Bailey, and Giuliano Maciocci. 2013. BeThere: 3D Mobile Collaboration with Spatial Input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 179–188.
31. Hyunyoung Song, François Guimbretière, Chang Hu, and Hod Lipson. 2006. ModelCraft: Capturing Freehand Annotations and Edits on Physical 3D Models. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST '06)*. ACM, New York, NY, USA, 13–22.
32. Yannick Thiel, Karan Singh, and Ravin Balakrishnan. 2011. Elasticurves: Exploiting stroke dynamics and inertia for the real-time neatening of sketched 2d curves. In *Proc. of UIST*. ACM, 383–392.
33. Anton van den Hengel, Anthony Dick, Thorsten Thormählen, Ben Ward, and Philip H. S. Torr. 2007. VideoTrace: Rapid Interactive Scene Modelling from Video. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07)*. New York, NY, USA, Article 86.
34. Vinayak, Devarajan Ramanujan, Cecil Piya, and Karthik Ramani. 2016. MobiSweep: Exploring Spatial Design Ideation Using a Smartphone As a Hand-held Reference Plane. In *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '16)*. ACM, New York, NY, USA, 12–20.
35. Shahtab Wahid, Stacy M. Branham, D. Scott McCrickard, and Steve Harrison. 2010. Investigating the Relationship Between Imagery and Rationale in Design. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems (DIS '10)*. ACM, New York, NY, USA, 75–84.
36. Christian Weichel, Manfred Lau, David Kim, Nicolas Villar, and Hans W. Gellersen. 2014a. MixFab: A Mixed-reality Environment for Personal Fabrication. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3855–3864.
37. Christian Weichel, Manfred Lau, David Kim, Nicolas Villar, and Hans W. Gellersen. 2014b. MixFab: a mixed-reality environment for personal fabrication. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 3855–3864.
38. 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)*. ACM, New York, NY, USA, 223–226.
39. Youyi Zheng, Han Liu, Julie Dorsey, and Niloy J Mitra. 2016. SmartCanvas: Context-inferred Interpretation of Sketches for Preparatory Design Studies. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 37–48.
40. Fabio Zünd, Mattia Ryffel, Stéphane Magnenat, Alessia Marra, Maurizio Nitti, Mubbasir Kapadia, Gioacchino Noris, Kenny Mitchell, Markus Gross, and Robert W. Sumner. 2015. Augmented Creativity: Bridging the Real and Virtual Worlds to Enhance Creative Play. In *SIGGRAPH Asia 2015 Mobile Graphics and Interactive Applications (SA '15)*. ACM, New York, NY, USA, Article 21, 7 pages.