shapeShift: 2D Spatial Manipulation and Self-Actuation of Tabletop Shape Displays for Tangible and Haptic Interaction

Alexa F. Siu, Eric J. Gonzalez, Shenli Yuan, Jason B. Ginsberg, Sean Follmer

Stanford University

Stanford, CA, USA

{afsui, ejgonz, shenli, jasang2, sfollmer}@stanford.edu

Figure 1. shapeShift is a mobile tabletop shape display that enables 2D spatial user input and output. *Left to right:* (a) In passive mode the user can freely move the display; (b) The display serves as a spatially-aware physical lens into underlying information; (c) Two displays can be used to simultaneously explore different regions; (d) In active mode shapeShift can move on-demand - here it tracks the user's hand to simulate the presence of virtual content

ABSTRACT

We explore interactions enabled by 2D spatial manipulation and self-actuation of a tabletop shape display. To explore these interactions, we developed shapeShift, a compact, high-resolution (7 mm pitch), mobile tabletop shape display. shapeShift can be mounted on passive rollers allowing for bimanual interaction where the user can freely manipulate the system while it renders spatially relevant content. shapeShift can also be mounted on an omnidirectional-robot to provide both vertical and lateral kinesthetic feedback, display moving objects, or act as an encountered-type haptic device for VR. We present a study on haptic search tasks comparing spatial manipulation of a shape display for egocentric exploration of a map versus exploration using a fixed display and a touch pad. Results show a 30% decrease in navigation path lengths, 24% decrease in task time, 15% decrease in mental demand and 29% decrease in frustration in favor of egocentric navigation.

ACM Classification Keywords

H.5.1 Multimedia Information Systems: Artificial, Augmented, and Virtual Realities; H.5.2 User Interfaces: Haptic I/O, input devices and strategies, interaction styles

Author Keywords

Shape-changing User Interfaces; Shape Displays; Actuated Tangibles

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

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

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-5620-6/18/04...\$15.00

DOI: <https://doi.org/10.1145/3173574.3173865>

INTRODUCTION

Our main ways of interacting with computers today are mostly centered around our visual sense and little advantage is taken of our inherent spatial reasoning skills and abilities to haptically interact with and manipulate the world around us. Actuated tangible user interfaces (TUIs) and devices that render haptic feedback have been proposed as methods to better leverage these innate skills. Among TUIs, shape displays are a type of input/output device that allow for more general purpose shape change. Work on shape displays has explored this spectrum through techniques for interacting with the surface [15, 48, 34, 26, 11], manipulating tangibles through the surface [11], and interacting at a distance through gestures [4]. We argue that this spatial spectrum [25] has mostly been constrained to the display's immediate vertical interaction space. A primary limitation of shape displays has been their large, heavy form factor and high manufacturing costs which have limited manufacturing of displays with large interaction areas. A large interaction area may be useful for exploration of virtual spaces and overcoming limitations in physical spatial resolution.

In this work, we attempt to overcome the workspace limitations of traditional shape displays to expand on the range of existing interactions enabled with them in rendering both shape content and UI elements through lateral 2D spatial input and output. Increasing the display size can provide a larger workspace with the tradeoffs of greater system complexity and cost. Thus, instead we propose *mobile shape displays* which use lateral mobility to both increase the workspace and allow for lateral motion as I/O. We introduce the concept of passive and active, or self-actuated, shape display mobility. In passive mode, a shape display is mounted on passive rollers such that its 2D movement on a tabletop is compliant and

entirely guided by user input. In active mode, a shape display is mounted on an omnidirectional robot platform that controls the movement of the display.

We leverage the passive and active mobile shape display platforms to demonstrate new interactions not previously capable with static shape displays. We explore new methods for bi-manual input, and show how the added translation and rotation degrees of freedom can be used as input to manipulate rendered objects. Using lateral self-actuation we show how shape displays are able to render objects' 2D movement over larger areas. We demonstrate applications of a mobile display as a physical lens for exploring spatial layouts (e.g., maps and volumetric data) and as an encountered-type haptic device for virtual reality (VR) scenarios using one or multiple displays.

To explore these interactions and applications, we introduce shapeShift: a high-resolution, compact, modular shape display consisting of 288 actuated pins ($4.85\text{ mm} \times 4.85\text{ mm}$, 2.8 mm inter-pin spacing) formed by six 2×24 pin modules. While shape displays remain limited in size and cost, shapeShift's self-contained hardware and small form factor allow it to be highly mobile on surfaces. shapeShift is capable of large lateral mobility, spatial continuity, and both vertical and horizontal kinesthetic haptic feedback.

In the last part of this work, we report on a user evaluation that explores the benefits of increased workspace and lateral I/O capabilities mobile shape displays afford. We hypothesized the passive mobility could be leveraged to help users better learn and understand spatial data by providing additional physical context through proprioception; thus, a spatial haptic search task was selected comparing static vs. mobile displays. As vision tends to dominate sensory input, we studied the effects of haptic and proprioceptive feedback alone. This is similar to the scenario in which one uses haptic feedback to explore or navigate an input device while visually attending elsewhere. Similar studies done for visual search tasks [7, 35, 29] have shown that egocentric body movements can improve recall of spatial positions; however the same has not been explored for encountered-type haptic devices and shape output.

The study on haptic search tasks compares egocentric spatial exploration through passive mobility of a shape display versus exploration using a static shape display and touch pad. Results show a 30% decrease in mean navigation path lengths, and 24% decrease in mean task time in favor of egocentric navigation. Moreover, self-reported task loads indicate a 15% decrease in mental demand and 29% decrease in frustration.

CONTRIBUTIONS

This paper provides the following core contributions:

1. Interactions with one or multiple mobile tabletop shape displays through passive movement and self-actuation.
2. Mobile shape displays as encountered-type haptics interfaces for dynamic rendering of virtual content in VR.
3. shapeShift, a new open-source modular hardware/software platform for a mobile tabletop shape display.
4. User study results on a haptic search task using a passively mobile tabletop shape display for egocentric exploration.

RELATED WORK

Shape-Changing & Surface Haptic Devices

Shape-changing user interfaces are a class of devices that use physical change in shape as input or output for human-computer interaction [37]. Shape displays are a type of shape-changing UI that enable rendering of objects and surfaces in 2.5D [15, 48, 34, 26, 11]. Project FEELEX [15] introduced malleable surfaces that combined haptic sensations with computer graphics, with the goal of enhancing visual content. The low resolution limited the type of content that could be represented. Follmer et. al. investigated how shape displays could be used for creating dynamic UIs [11]; however, they were not able to explore UI elements implementing lateral kinesthetic haptic feedback.

Robotic Graphics & Encountered-Type Haptics

McNeely introduced the concept of robotic graphics as a new approach for providing force feedback in virtual environments, where the user physically interacts with external robots simulating a virtual object [28]. Hirota et al. furthered this concept with their Surface Display, a 4-by-4 pin array mounted to a passive gimbal used to render surfaces of virtual objects [14]. More recent work has explored the use of a robotic arm with end effector-mounted surfaces to provide haptic feedback [50, 2]. shapeShift builds upon previous work by enabling dynamic rendering of virtual surfaces presented to the user.

Self-Actuated Objects on Interactive Surfaces

Typically tangible user interfaces (TUIs) have been used as input controls [49], for manipulating remote physical environments [31, 38, 39] and to display information [9, 47, 27]. Recently, tabletop swarm user interfaces such as Zooids [22] have been used as a type of input/output interface where some agents act as controls or handles while others are used for shape output. They have also been used for assembly of on-demand objects in VR [51]. Compared to work on shape displays, they are capable of movement in a 2D plane to provide tangential feedback and have a much lower fabrication cost. With shapeShift we aimed at creating a mobile tabletop display to exploit some of the benefits from both.

Egocentric Navigation with Spatially-Aware Devices

Previous work has shown that having stable information spaces and spatial visual cues improves user performance of desktop UI navigation since users quickly learn locations as a side effect of use [17, 8, 24]. Traditional desktop interfaces have mostly exploited using visuals to provide users with spatial cues. However, sound [21] and proprioception through egocentric body movements [19, 7] even without vision can also provide users with information about spatial relationships.

Spatially-aware devices can be more effective for line-length discrimination [29], map navigation [40, 35, 36], and 3D volume exploration [46, 20]. These systems allow users to visualize data in slices spatially located instead of fixed onto a two dimensional screen. Hover Pad proposed a similar system but wherein the display moved autonomously to overcome limitations from the user's physical input [42]. shapeShift extends the idea of situated information spaces to enable physical exploration of data.

INTERACTIONS WITH MOBILE SHAPE DISPLAYS

Input Through Movement

With a mobile tabletop shape display, users have control of an additional 3 degrees of freedom: translation in the x - y plane and rotation about the z -axis. Control of these additional degrees of freedom can be thought of from two perspectives.

Figure 2. A mobile shape display can be moved over content using (a) the viewport itself or (b) objects or handles fixed relative to the viewport.

In the first case, movement of the display translates to movement of a viewport into spatially-situated underlying information. Movement of the display changes the viewport's position relative to the underlying information, changing the view (Figure 2a). This concept is similar to *Situated Information Spaces* [10] and *Tangible Lens* [25]. The display platform may have additional side handles that allow the user to easily move it.

In the second case, the user is spatially manipulating rendered content fixed relative to the display. Moving the content moves the display, as the two are coupled. As the content is moved in space, other content around it is hidden/revealed (Figure 2b).

With multiple objects in the space, users may want to select which object they want to manipulate. Various methods can be used to know which object the user is *selecting* to manipulate; for example, one could define selection gestures or detect the object the user is touching through hand tracking. The user might also want the freedom to switch between object selection and viewport movement. Touching the sides of the display could indicate viewport movement while directly touching the pins could indicate object selection.

Content movement can also be used to define UI elements. We illustrate the use of physically rendered handles to translate/rotate a shape display in Figure 3. Elements can be pushed or pulled in the x - y plane or rotated about the z -axis for input.

Bimanual & Unimanual Interaction

Compared to static displays, a mobile tabletop shape display with its additional degrees of freedom lends itself to two-handed, or bimanual interaction. We outline a range of scenarios in Figure 4.

Figure 3. Example lateral UI elements for a passively mobile shape display: (a) slider, (b) joystick, and (c) knob

Figure 4. Bimanual scenarios for a mobile tabletop shape display.

In the simplest case of a static display, one or two hands may be used to interact with the content (Figure 4). With translation and rotation involved, positioning and interaction with the content could be more efficient as an asymmetric compound task. The non-dominant hand performs coarse positioning of the display while the dominant hand provides support for fine positioning (e.g., control for rotation). Once the desired location is reached, the non-dominant hand provides support to prevent further movement, while the dominant hand is free to explore the display content.

Past work has shown the benefits of using two hands for compound tasks [6], two-handed coordination is most natural when the dominant hand moves relative to the non-dominant hand [13]. Hinckley et. al. showed in a virtual object manipulation experiment that when using two hands as opposed to unimanual control, participants were able to maintain a fairly precise body-relative representation of space [12]. We further explore this idea in our user evaluation.

Output Through Self-Actuation

Active mobile shape displays can be used to display 2D spatial movement of objects. Actuation of the pins can render the object's physical form, while self-actuation of the platform can show its movement in space (Figure 5).

Figure 5. Shape display self-actuation can be used to display spatial movement of objects. Here we illustrate the concept with a bouncing ball. Actuation of the pins renders the ball's vertical movement while actuation of the platform shows its projectile motion in space.

Self-actuation can also be used for kinesthetic haptic feedback (Figure 6). The vertical movement of the pins provides vertical kinesthetic feedback while self-actuation of the display provides horizontal kinesthetic feedback. Objects can resist movement from users or they can react and push back on users hands. As such, display elements can be used to create position or movement constraints. This is especially useful for rendering UI elements, such as those in Figure 3.

Interaction With Multiple Mobile Shape Displays

Multiple mobile tabletop shape displays increase the degrees of freedom for interaction, and can be joined to increase the interaction space on a 2D plane and to render larger objects.

Alternatively, with separate mobile shape displays for each hand, a user may simultaneously explore multiple locations

Figure 6. An active mobile shape display can be used to provide haptic feedback. Actuation of the pins provides vertical kinesthetic feedback (red arrows) while self-actuation of the platform provides horizontal kinesthetic feedback (blue arrows).

of the underlying content (Figure 1c). Compared to multiple aggregated displays, this further increases the interaction space but limits interaction to a single user. Here the mobile display movement may be active or passive. In the active case, a self-actuated platform and 3D tracking are needed to track the user's hands and render the content beneath them. We discuss an implementation of the active case for VR applications in the next section. In the passive case, movement of the viewports are directly controlled by the user's hands, which may be free or coupled to the displays.

Two displays can also be used to simulate an *infinite* surface (Figure 7). With one display, when the user reaches the end of the display margin, no more content can be haptically explored. With an additional self-actuated mobile display, the second display can move into place to continue rendering the next segment of the content. The two displays, can continue to move synchronously like in a relay race, so the user never feels a discontinuity in their haptic exploration. This would require knowledge or accurate prediction of where the user will move their hand next.

Figure 7. Multiple active tabletop displays can be used synchronously to simulate an infinite surface.

Application Demonstrations

Rapid Physical Rendering and Manipulation of 3D Models. Using a passive mobile shape display, rendered objects can be freely grasped, translated, and rotated allowing for more natural manipulation. This is particularly useful to designers reviewing physical models or layouts, who often need to quickly change viewpoints and access different areas of a design. In Figure 8 (left), we show shapeShift rendering a game controller and the user rotating it to explore different features.

Physical Lens for Exploring Spatial and Volumetric Data. A mobile shape display can be moved along a surface to explore spatial information that extends beyond the boundaries of the display. It can be used as a stand-alone physical display or to enhance a graphical user interface, such as overhead projection or augmented reality. Movement can be passive, where the

Figure 8. In a passive mobile display, rendered objects can be freely grasped, translated and rotated (left) spatial data such as terrain maps can be mapped to real-world spaces for physical exploration (right).

Figure 9. Virtual content such as a terrain map can be modified and physically explored real-time using physical proxies such as a wand (left). Two displays are used to physically render two different virtual houses (right).

user positions the display in a region of interest where tangible detail is desired, or it can be active, where the display moves itself to guide the user's attention or highlight a region. For example, if exploring seawater levels on a geographical map, data filters can be applied to render the region with highest water level. We show shapeShift in Figure 1b, rendering a region of a torus. As the user moves the display over the table, different parts of the torus are rendered. In Figure 8 (right), we use shapeShift to physically render a terrain map.

ENCOUNTERED-TYPE HAPTICS FOR VR

An active tabletop shape display can be used as an encountered-type haptics display for VR applications. Combined with overhead IR tracking cameras and a glove fitted with IR-reflective markers, the system can track the user's hand and physically render the content the user is interacting with in the virtual world (Figure 1d). Kinesthetic feedback can be provided vertically by the actuation of the pins or laterally by movement of the robot platform. Compared to previous work [50, 2], pin actuation allows us to render dynamic surfaces rather than a single static surface.

Users can reach into the virtual scene to feel, push, and grab virtual objects (Figure 1d), allowing users to interact with virtual world objects more intuitively and efficiently. This can be combined with other physical proxies, such as hand-held controllers, that can be used to modify the virtual content rendered by shapeShift.

In addition, shapeShift can also be used to reach out to the user, such as an object nudging the user to direct their attention or presenting a new surface to them.

VR Application Demonstrations

We used shapeShift to physically render a terrain map in VR (Figure 9b). Using an IR-tracked physical wand, the user could raise and lower the terrain. This allows them to visualize and feel changes made to the virtual terrain in real-time.

Figure 10. shapeShift renders physical shapes using 288 actuated square pins (4.85 mm width, 7 mm pitch). Passive and active mobile tabletop platforms enable lateral motion, with user interaction and platform motion both tracked using either an IR motion capture system (shown) or an HTC Vive-based tracking system.

Two shapeShift devices were used together to enable two-handed free exploration of a spatial map (Figure 1c). Multiple displays can also be used to each represent individual objects in a shared virtual space. In Figure 9a, we show two versions of shapeShift, where each one represents a different house (house A and house B). The user can translate and rotate each house to position them in the virtual world.

SYSTEM IMPLEMENTATION

Design Rationale

Size of Rendering Region. Since haptic shape perception is often done with our hands, a hand-sized display may be sufficient for initial exploration of the design space. Thus the minimum XY display size (9cm x 17.5cm) was selected to be close to the mean maximum human hand spread (21cm) [16] with the assumption that maximal spread is rare and that a single display is meant for a single hand; however, we also chose a modular design to allow expansion as needed. The electronics are mounted between rows, forming modules that can be stacked to expand the rendering region.

Spatial Resolution. In selecting spatial resolution, the two-point discrimination threshold of 1.25 mm would be ideal [5]. In practice, optimizing around actuator size, cost, and modularity yielded shapeShift's pin pitch of 7 mm. To reduce complexity and overall display height, the motors were aligned 1:1 with the pins; thus, the display resolution was limited by actuator diameter (TTMotors TGPP06, 1:25, 6 mm). These actuators were selected not only for their small size but also to keep the individual price at a reasonable cost (<3.50 USD/pc).

Mobility. To enable different forms of lateral interaction, the display needed to be light enough to be easily manipulated with a single hand. Weighing 6.8 kg, shapeShift can be swapped between passive and active mobile platforms based on the desired application. Passive movement of the platform may cause user fatigue over time, thus we explore the self-actuated platform as an alternative.

Physical Rendering Latency. The physical speed at which the pins can move limits the speed at which objects can translate or appear on the surface. Thus, there is a maximum speed $V_{display,minLag}$ with which the display can travel (or in the case of a static display, that with which content moves across the display) to ensure minimal lag between lateral displacement and vertical pin rendering. This speed limit is quantified by

$$V_{display,minLag} = V_{pin,max} * |(\nabla S \cdot \hat{u})^{-1}| \quad (1)$$

where $V_{pin,max}$ is the maximum pin speed, S is the surface to be rendered, and $\hat{u}$ is the unit vector in the display's (or content's) lateral direction of travel. Assuming no software latency, speeds below this limit would produce no physical lag. For example, in order to ideally render a plane of slope 1 mm/mm shapeShift must move at or below 75 mm/s, based on a max vertical pin speed of 75 mm/s; at speeds faster than this, the pins take longer to render the surface than the display takes to move over it. In theory, if the maximum surface gradient in any given direction was known, the speed limit for any object's ideal physical rendering could be calculated. In practice, discontinuities, sharp edges, and software latency prevent rendering with no physical lag, although perceived lag could be considerably reduced by increasing pin speed.

Hardware

Based on the design rationale we created ShapeShift. A system schematic is shown in Figure 10. The display consists of 288 hollow aluminum pins (4.85 mm x 4.85 mm) arranged in a 12 x 24 grid with an inter-pin spacing of 2.8 mm. Each pin is 152 mm in length and can travel up to 50 mm. The overall rendering volume measures 178 mm x 89 mm x 50 mm.

Each pin is coupled to a nut and individually actuated by a custom 3 mm pitch motorized lead screw, powered by a DC Motor (TTMotors TGPP06, 1:25, 6 mm). Guide grids keep the pins aligned and prevent rotation, which results in linear motion. To sense the pin's position, two photointerruptors detect quadrature tick marks on a 3D printed shaft coupling between the motor and screw; tick counts are then integrated to obtain a reading proportional to linear position. A limit switch (Omron D3SH-A0L) at the base of the screw is used to calibrate all pins at program startup. The average position accuracy per pin is 0.35 ± 0.22 mm. The average pin speed is 75 mm/s traveling upwards and 81 mm/s downwards.

Each pair of adjacent rows functions as a module of 2×24 pins (Fig 11a), independently controlled by two double-sided PCBs mounted side-by-side between the two rows. While modules can be stacked as desired in the x-dimension to expand the rendering region (Fig 11b), the current implementation uses 6 modules or 12 PCBs total. A custom motor mount bridges the two PCBs, accommodating the 24 motors per row.

Figure 11. shapeShift is made of stackable 2×24 pin modules (a), that can be combined to expand the display size in one dimension (b).

Figure 12. Block diagram of the shapeShift system. Motion is captured and processed in Unity, where a script calculates the pin positions for rendering content at the given position/orientation. A low-level control loop governs the position of each of the n pins, where each MCU manages 6 pins. In this implementation, n is 288.

Each PCB has four ARM Cortex-M4 72 MHz microcontrollers (NXP, MK20DX256VLH7) each managing 6 pins via PID control at 500 Hz. All MCUs communicate in serial over a 10^6 bps RS485 shared bus (MAX 3465). An external MCU serves as master for the display, communicating via USB serial with the computer and forwarding messages to the slaves via RS485. An overall block diagram is shown in Figure 12.

The system consumes near 140 W (23 W per module) on average supplied by two 5 V 60 A power supplies. In the worst case of all motors being stalled, based on a 1 W peak power consumption per motor, the display could consume up to 336 W. Heat dissipation is managed by 4 PC cooling fans.

Mobile Platform

In passive mode, the display is mounted to a caster wheel platform, allowing the user to freely move the display. In active mode, the display is mounted to an omnidirectional mobile robot (Nexus Robot 10008) and the feedback loop is closed with position tracking data. An ATmega328 microcontroller controls the platform, which has a max speed of 600 mm/s. The current system measures $26.7 \text{ cm} \times 24.4 \text{ cm} \times h$ in, where h is 29.5 cm in passive mode and 38.4 cm in active mode.

Figure 13. A ray casting script is used to check collisions with objects in a virtual world and determine the target position for each pin in the display. If a collision does not occur, the pin is set to its initial position.

Position Tracking

To track the display as it moves in the environment, we use one of two methods. In the first method, we use a system of overhead OptiTrack cameras¹ and IR-reflective markers to track the display's position and orientation. Alternatively, we have also implemented 3D tracking using the HTC Vive Lighthouse system and a ViveTracker² mounted on the display.

Software

A Unity application is used to import 3D models, interface with the tracking device (OptiTrack or ViveTrackers), and calculate target pin positions depending on what is being rendered. The application runs two additional parallel threads to handle USB serial communication. One thread communicates with the master microcontroller any updates to pin position or display settings (e.g., PID parameters). The second thread communicates with the omnidirectional robot when the display is used in active mode.

To determine target pin positions, the application simulates a plane the size of the 2D rendering region (Figure 13). At each pin location, a ray is cast into the virtual world from a known height, H . If the ray collides with a virtual object at height δ , the pin position, p , is set to: $|p| = |H - \delta|$. If no collision is detected, then the pin position is set to its initial position.

The data transmission pipeline for the system given 3D position data from an arbitrary tracking system is as follows: Calculate target pin positions via ray casting (0.20 ms/module) → Send positions to Master MCU via USB serial (0.04 ms/module) → Forward positions to display MCUs via RS485 (1.63 ms/module).

For the current implementation of 6 modules (288 pins) the total data transmission latency is 11.2 ms. Overall, we implement a 50 Hz data refresh rate for the display.

USER EVALUATION

One of the main contributions of this paper is spatial manipulation of a mobile shape display for exploring spatial data. Our hypothesis is that a passively mobile shape display will perform better than interactions with a shape display that is fixed in location (static) since the additional body movements provide additional physical context through proprioception. To test this, we designed a task where participants used shapeShift

¹<http://optitrack.com/products/prime-13/>

²<https://www.vive.com/us/>

Figure 14. Primitive shapes used in the object recognition task (1-10) and spatial navigation task (1-6). The last row shows examples of how the shapes (1-3) looked when rendered by shapeShift.

to haptically explore a map and find given target shapes. Participants were only able to feel the map features as they were rendered by shapeShift. The only visual cue that was given was a screen which showed real-time position of the viewport relative to the map but did not reveal any of the map features. We removed visual cues, so it would not dominate over haptic and proprioceptive cues.

Similar experiments have been performed using 2D screens for egocentric navigation [35, 40, 29, 18]; in these scenarios, participants have visual feedback of the map. For haptic search tasks within the peripersonal or manipulatory space [23], it has been shown that proprioception allows participants to return their hand to previously touched target locations with low error rates [19, 30, 33]. However, tests on long-term spatial memory recall from map navigation have not been explored.

Selecting Map Features

Shimojo et. al. conducted shape recognition studies in a tactile display by varying pin pitch [43]. They found a 2 mm pitch to be ideal for recognition of shapes between 30 mm to 48 mm diameter with a single finger. While shapeShift does not meet these specifications, in our own experience, it can render shapes which can be recognized. Therefore we needed to conduct a study to understand which shapes could be easily perceived. These shapes would then be used for the spatial map. We conducted a short object recognition study where participants identified shapes without visual cues; a blocking screen prevented participants from seeing the rendered shapes.

We recruited 8 participants (3 female, mean age 25.8, std. dev. 2.0). Participation was voluntary and the entire experiment lasted approximately 8 minutes.

We created a dataset of ten different shape primitives from *surface type*, *edge type*, and *vertex type* shape categories similar to those used in the experiments by Shimojo et. al. and Sinclair et. al. [43, 44]. These shapes are shown in Figure 14. The footprint of each rendered shape was 75 mm × 75 mm, except the hemiellipsoid which had a minor axis of 50 mm. The maximum height of each shape was 50 mm. Compared to the experiments of Shimojo et. al., participants' exploratory movements were left unconstrained, and the shapes used are twice as large (span the entire display width).

Object recognition results are summarized in a confusion matrix on Figure 15. Participants took on average 6.1 s to identify

Figure 15. Confusion matrix summarizing results from an object recognition study ($N = 8$).

Figure 16. a) Shows the user study setup. b) Shows a participant using shapeShift in the passively mobile condition. c) Shows a participant using shapeShift in the static condition.

each shape and were generally successful with 30% of shapes having 100% recognition success and the lowest recognition success being 87.5%. From these results, we identified shapes 1-6 as the most unique to be used in the haptic search task.

Haptic Search Study Conditions & Setup

The general setup is shown in Figure 16a. There were two conditions for the haptic search task: static condition and passive condition.

In the passive condition, participants explored the map by physically moving the shape display on a tabletop (Figure 16c). Movement of the viewport was coupled to shapeShift's physical movement. shapeShift's position was tracked using a ViveTracker, and a virtual map was situated within the table boundaries so users only had to move shapeShift within the boundaries to explore the map.

In the static condition, the display remained fixed in place on the table (Figure 16b). To explore the map, participants moved the virtual viewport position using a trackpad.

For both conditions, navigation was performed using bimanual input. The non-dominant hand was used to move the viewport while the dominant hand identified the map features. The dependent variables were navigation time, path length, map memory recall score and task load questionnaire responses.

Materials

Maps used for the haptic search task measured 0.43 m x 0.80 m. This workspace size was chosen to be within the reported average size for zone of convenient reach (ZCR) based on anthropometric data [32, 41].

Participants explored two maps for each condition. Each map had six features (Figure 14, 1-6); four were randomly chosen as targets while the remaining two were used as distractors. The features were randomly arranged in the map such that the distance between targets was no more than 0.2 m. The targets were repeated twice, such that for each map, participants had to find the target exactly two times. Therefore for each condition there were a total of 2 maps × 8 targets/map = 16 trials.

For the static condition, the trackpad used was an Apple Magic Trackpad 2 (6.30 in × 4.52 in). The trackpad sensitivity was set such that the average speed of the viewport in both conditions was approximately the same, as determined from pilot testing.

For both conditions, participants had view of a small screen (7 in, 1024 × 600) that showed the location of the display within the workspace limits but had no information about the targets or distractors (Figure 16a).

Participants

We recruited 13 right-handed participants (7 female, mean age 28, std. dev. 6.4). Participants received 15 USD compensation and the experiment lasted approximately 60 minutes.

Procedure

The experiment began with a training map that was not the same as the maps used for the experimental condition. Participants were asked to locate targets until they felt comfortable with the different input methods. The targets used in the practice training map (i.e. a bar, a capsule, and oval) were different from those in the actual study.

A trial began when participants verbally acknowledged being ready. A target name would then be revealed on the screen and participants were told they could immediately start exploring the map. They were told their performance would be measured based on *accuracy*, *time taken to reach the target*, and *length of path traveled to reach the target*. If after three minutes the target had not been found, the trial was terminated. All participants found the targets within the time limit.

Participants were told to verbally acknowledge when they were confident the target had been found and stop moving the display. Their answer was recorded and time was logged. The next trial began when they again acknowledged being ready, with the display starting from where the previous trial ended.

After each map (i.e., completion of eight trials), participants were given a semantic memory distraction task for two minutes to control for recency effects [3, 45]. During this task participants were asked to list as many US states, cities in

Path Length, Time & Speed vs. Trial Number

Figure 17. Mean path length, task time, and speed for all trials and participants ($N = 13$) versus trial number. With more trials, navigation path and time decrease.

California, or countries in Europe or Asia as they could within the time limit.

Afterwards participants were asked to draw the map they had just explored to the best of their ability and complete a memory recall questionnaire that asked them questions about the spatial arrangement of features relative to each other in the map. For example, "Was *cylinder* to the right or left of *hemisphere*?" For each answer, users also had to rate on a scale from 1 to 5 how confident they were in their response. For each question, a weighted score was computed by multiplying the score times the confidence; this yields values between 0 to 5.

At the end of each condition, participants completed a NASA-TLX questionnaire followed by a two minute break before continuing with the next condition. Order of conditions was counterbalanced between subjects. After completion of both experimental conditions, participants were asked to complete a post-survey with questions regarding their overall experience and input preferences.

Hypothesis

We hypothesized that participants would be able to more accurately recall the maps in the passive condition than static condition since having to move their hand would provide additional proprioceptive cues (H1). As a result, we also expected, participants would take less time and shorter paths for each trial in the passive condition (H2). We also hypothesized that while the passive condition could result in higher physical demand, the mental demand would be lower (H3).

Results

To assess a user's learned mental map, we compared task completion results between early trials when participants had not yet seen all targets (the learning phase) to later trials when participants had seen each target at least once (the navigation phase). This approach is mirrored from literature on spatial learning using peephole navigation [35].

Figure 18. Mean path length, task time, and speed after participants had explored the targets at least once ($N = 13$). Mean navigation path length decreased by 30% and the task time decreased by 25% for the passive condition. No significant differences were found in speed of movement.

Path, Time, and Speed Before Exploring All Targets

Three two-tailed dependent t-tests with 95% confidence interval and Bonferroni-Holm correction were used to determine whether there was a statistically significant mean difference in path lengths, time, and speed between conditions for the initial trials (1-4) when participants had not yet explored all targets. The assumption of normality was not violated, as assessed by a Shapiro-Wilks test ($p > 0.05$ for all three tests). No significant differences were found in mean path length ($t(12) = 0.207, p = 0.83$), time ($t(12) = 0.55, p = 0.59$), and speed ($t(12) = 1.57, p = 0.14$) between conditions for the first four trials. These results show that both input methods are comparable for the task. Figure 17 shows the mean path length, time, and speed plotted versus trial number for all participants.

Path, Time, and Speed After Exploring Targets At Least Once

Three two-tailed dependent t-tests with 95% confidence interval and Bonferroni-Holm correction were used to determine whether there was a statistically significant mean difference in path lengths, time, and speed between conditions for the last trials when participants had seen each target at least once. The assumption of normality was not violated, as assessed by a Shapiro-Wilks test ($p > 0.05$ for all three tests). Results are summarized in Figure 18. Significant differences were found in mean path length ($t(12) = -2.85, p = 0.014$) and time ($t(12) = -2.95, p = 0.012$) between conditions. No significant differences were found for speed ($t(12) = 0.39, p = 0.70$). In the passive condition, the average navigation path length decreased by 30% and the average task time decreased by 25% when compared to the static condition.

Memory Recall Questionnaire

A fourth t-test with 95% confidence interval and Bonferroni-Holm correction was used to determine whether there was a statistically significant mean difference in weighted recall score between the passive ($\mu = 3.47, \sigma = 0.78$) and static conditions ($\mu = 3.02, \sigma = 0.92$). No significant differences were found ($t(12) = 1.98, p = 0.07$).

NASA-TLX Questionnaire

Results from pairwise comparisons between subscales in the NASA TLX questionnaire are summarized in Figure 19. Significance values were determined based on two-tailed dependent t-tests. Significant differences were found in mean mental demand ($t(12) = -2.54, p = 0.02$) and frustration

Figure 19. Self-reported mean task load ratings for the study task in both conditions ($N = 13$). Mean mental demand decreased by 15% and mean frustration decreased by 29% in favor of the passive condition.

($t(12) = -2.56, p = 0.02$) between conditions. There was a 15% decrease in mental demand and 29% decrease in frustration in favor of the passive condition.

Qualitative Feedback

Twelve out of the fourteen participants chose the passive shape display on wheels as a more intuitive interface for moving the viewport and successfully accomplishing the study task. The passive movement helped by giving them a better sense of location, "Felt easier to position accurately and gave me a better sense of where everything was since I was physically moving." They felt the physical movement was better suited for the task, "It felt good to move TO the object. The track pad, on the other hand, felt like bringing the object to me."

Reasons for finding the trackpad less intuitive include, "...I had to imagine that I was moving, rather than actually moving.", and "Had to rely on visual display a lot more to tell where I was. Lost track often when I want to go quick."

From the two participants that chose the trackpad as a more intuitive input, the main complaints were related to disliking having to also control the rotation degree-of-freedom of the passive shape display, "I had to use more effort to move and position the device properly. It is sometimes hard to keep the device parallel (prevent from rotation)".

Discussion

Overall results show the passive tabletop display was more favorable for the spatial haptic search task. Lower mean navigation path and time show that participants were able to perform the task more efficiently. Speed of movement was not an issue since difference in speed between conditions was statistically insignificant. The speed in both conditions may have been influenced by the system's physical rendering latency. Taking the pyramid as an example, with a max planar slope of 1.33 mm/mm the speed limit for its minimal lag rendering is 56 mm/s according to Equation (1). This closely matches the mean travel speed in both conditions (Figure 18), indicating that user speed may have been determined by the rendering latency of the system.

Moreover, when analyzing the first trials where participants had not yet seen the targets, no differences were found between conditions in mean navigation path, time, or speed. With more trials, however, results show the average path length and time starts to decrease more quickly for the passive condition (Figure 17). This could mean participants were able to more quickly learn the map relationships.

The greater task efficiency could be explained by the lower cognitive load the passive display has over the static display. Users reported on average lower mental demand and frustration when using the passive display. By having more spatial persistence in the passive condition, users may have been able to leverage use of inherent spatial reasoning skills and muscle memory. This is reflected in some of the users comments, where in the trackpad condition they felt restricted by being unable to simply reach out and find the object.

The lower cognitive load could also be attributed to lower dependency on the visual position cues. Integrating visual cues with the trackpad hand movement requires an additional cognitive step compared to bimanual movement. As shown by Hinckley et. al., with two hands users can maintain a precise, body-relative spatial representation which is not dependent on visual feedback. [13]. In a trajectory-based device evaluation with visual feedback [1], Accot et. al. found an absolute mapping device (e.g., tablet + stylus) increased task performance compared to relative mapping devices (e.g., touchpad) due to the higher dexterity afforded by the stylus and removing the need to "clutch" during long travel segments. The same trend presumably applies to a spatial haptic search task.

We were not able to prove our hypothesis that spatial persistence and proprioceptive cues from egocentric movement could help users in long-term memory recall. These results are similar to [35]. While their study involved a visual search task and not a haptic search task, their results showed significant differences in path length and time but no differences in long-term memory recall. In their study, differences in recall were only observed in a follow-up study where the distraction task was 15 minutes long. For our study, while the mean recall score for the passive condition was higher than that of the static condition, the difference was not significant. Perhaps a longer distraction task would be necessary in order to observe the effects on long-term memory recall.

LIMITATIONS & FUTURE WORK

A significant limitation of any interactive shape display is the latency of physical rendering, but even more so for spatially-aware mobile displays which constantly reveal/hide content as they move in space. While data latency may be improved by segmenting the display into separate data buses, the main bottleneck of shapeShift is pin speed. This is due to our choice of actuators, which were selected for low cost, high-density packing at the cost of speed. While shapeShift's current rendering speed is fast enough for many real-time interaction scenarios, for scenarios involving more rapid movement it may not be. When using a shape display as an encountered-type haptic device for VR, it may be possible to use visuo-haptic illusions to mitigate these speed limitations.

VR applications using self-actuated tabletop shape displays may also benefit from more sophisticated hand-tracking techniques. In this work, we only investigate the approach of exactly following a user's hand; however, estimation of hand trajectory and touch prediction could further improve VR interactions by enabling shape displays to preemptively move to/render objects the user is expected to touch. Rendering could be further improved by detecting the shape display's optimal orientation to render an object to minimize discrepancy between the virtual and physical surface (e.g., the display self-rotates so its grid aligns with the primary axes of a cube).

Without relative motion between the hand and display (e.g., during hand-following), the pins can only exert a normal force on the hand as content is rendered, not shear force. Further study is required to investigate the haptic perception of relative motion between the hand and self-actuated mobile displays.

Currently, direct touch input is not implemented in shapeShift; information about where/what the user is touching is limited to spatial tracking data. Pin-level touch and force sensing (e.g., via capacitive sensing/load cells) would expand shapeShift's capabilities as both a tangible UI and haptic device.

CONCLUSION

Mobile shape displays combine the benefits of both actuated tangibles (dynamic UIs with controllable lateral kinesthetic haptic feedback) and actuated pin arrays (general purpose shape-change). Mobility allows us to leverage *smaller* shape displays to provide users with a large workspace while maintaining lower costs and complexity. Through shapeShift, we have explored the passive and active mobility of shape displays. This type of display can be used to render and manipulate static content or to physically explore 2D situated information spaces. Our user evaluation on a spatial map search task using a passively mobile shape display as navigation input shows that it helps increase search efficiency and reduce users' cognitive load.

Self-actuation of tabletop shape displays enables new forms of output such as displaying objects' 2D spatial motion and providing lateral kinesthetic feedback. Combined with an HMD, the display can also be used as an encountered-type haptic device for VR applications.

The added degrees of freedom of mobile tabletop shape displays offer an exciting area of opportunity to better leverage our haptic senses and spatial skills when interacting with digital information. We hope shapeShift's open-source platform will inspire both designers and researchers in prototyping new interactive physical interfaces.

All necessary documentation for implementing shapeShift can be found at <https://github.com/ShapeLab/shapeShift.git>.

ACKNOWLEDGEMENTS

We thank Mathieu Le Goc and Allen Zhao for their advice and support. This work was supported in part by the NSF GRFP under Grant No. DGE-114747, Stanford University School of Engineering Fellowship, Volkswagen, and the Hasso Plattner Design Thinking Research Program.

REFERENCES

1. Johnny Accot and Shumin Zhai. 1999. Performance evaluation of input devices in trajectory-based tasks: an application of the steering law. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 466–472.
2. Bruno Araujo, Ricardo Jota, Varun Perumal, Jia Xian Yao, Karan Singh, and Daniel Wigdor. 2016. Snake Charmer: Physically Enabling Virtual Objects. In *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 218–226.
3. Robert A Bjork and William B Whitten. 1974. Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology* 6, 2 (1974), 173–189.
4. Matthew Blackshaw, Anthony DeVincenzi, David Lakatos, Daniel Leithinger, and Hiroshi Ishii. 2011. Recompose: direct and gestural interaction with an actuated surface. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1237–1242.
5. Patrick Bruns, Carlos J. Camargo, Humberto Campanella, Jaime Esteve, Hubert R. Dinse, and Brigitte RÄüder. 2014. Tactile Acuity Charts: A Reliable Measure of Spatial Acuity. *PLOS ONE* 9, 2 (02 2014), 1–11. DOI: <http://dx.doi.org/10.1371/journal.pone.0087384>
6. W Buxton. 1995. Touch, Gesture & Marking. Chapter 7 in RM Baecker, J. Grudin, W. Buxton and S. Greenberg, S.(Eds.). *Readings in Human Computer Interaction: Toward the Year 2000*. (1995).
7. Emanuele Coluccia, Irene C Mammarella, Rossana De Beni, Miriam Ittyerah, and Cesare Cornoldi. 2016. Remembering Object Position in the Absence of Vision: Egocentric, Allocentric, and Egocentric Decentered Frames of Reference. *Perception* 36, 6 (June 2016), 850–864.
8. Brian D Ehret. 2002. Learning where to look: Location learning in graphical user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 211–218.
9. Katherine M Everitt, Scott R Klemmer, Robert Lee, and James A Landay. 2003. Two worlds apart: bridging the gap between physical and virtual media for distributed design collaboration. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 553–560.
10. George W Fitzmaurice. 1993. Situated information spaces and spatially aware palmtop computers. *Commun. ACM* 36, 7 (1993), 39–49.
11. Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. 2013. inFORM: Dynamic Physical Affordances and Constraints Through Shape and Object Actuation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, 417–426. DOI: <http://dx.doi.org/10.1145/2501988.2502032>
12. Ken Hinckley, Randy Pausch, and Dennis Proffitt. 1997a. *Attention and visual feedback: the bimanual frame of reference*. ACM, New York, New York, USA.
13. Ken Hinckley, Randy Pausch, Dennis Proffitt, James Patten, and Neal Kassell. 1997b. Cooperative bimanual action. In *the SIGCHI conference*. ACM Press, New York, New York, USA, 27–34.
14. Koichi Hirota and Michitaka Hirose. 1995. Simulation and presentation of curved surface in virtual reality environment through surface display. In *Virtual Reality Annual International Symposium, 1995. Proceedings*. IEEE, 211–216.
15. Hiroo Iwata, Hiroaki Yano, Fumitaka Nakaizumi, and Ryo Kawamura. 2001. Project FEELEX: adding haptic surface to graphics. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 469–476.
16. Garrett JW. 1971. The adult human hand: some anthropometric and biomechanical considerations. *Human Factors* 13, 2 (April 1971), 117–131. DOI: <http://dx.doi.org/10.1177/001872087101300204>
17. Victor Kaptelinin. 1993. Item recognition in menu selection: the effect of practice. In *INTERACT'93 and CHI'93 conference companion on Human factors in computing systems*. ACM, 183–184.
18. Frederic Kerber, Antonio Krüger, and Markus Löchtefeld. 2014. Investigating the effectiveness of peephole interaction for smartwatches in a map navigation task. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. ACM, 291–294.
19. R L Klatzky and S J Lederman. How well can we encode spatial layout from sparse kinesthetic contact?. In *11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. IEEE Comput. Soc, 179–186.
20. Jonathan Konieczny, Clement Shimizu, Gary Meyer, and D'nardo Colucci. 2005. A handheld flexible display system. In *Visualization, 2005. VIS 05. IEEE*. IEEE, 591–597.
21. James R Lackner and Barbara Shenker. 1985. Proprioceptive influences on auditory and visual spatial localization. *Journal of Neuroscience* 5, 3 (1985), 579–583.
22. Mathieu Le Goc, Lawrence H. Kim, Ali Parsaei, Jean-Daniel Fekete, Pierre Dragicevic, and Sean Follmer. 2016. Zooids: Building Blocks for Swarm User Interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, 97–109. DOI: <http://dx.doi.org/10.1145/2984511.2984547>
23. Susan J Lederman, Roberta L Klatzky, April Collins, and Jackie Wardell. 1987. Exploring environments by hand or foot: Time-based heuristics for encoding distance in movement space. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13, 4 (1987), 606.

24. Eric Lee and James Macgregor. 2016. Minimizing User Search Time in Menu Retrieval Systems. *Human Factors* 27, 2 (Nov. 2016), 157–162.
25. Daniel Leithinger. 2015. *Grasping information and collaborating through shape displays*. Ph.D. Dissertation. Massachusetts Institute of Technology.
26. Daniel Leithinger and Sean Follmer. 2010. *Relief: a scalable actuated shape display*. ACM, New York, New York, USA.
27. Daniel Leithinger, Sean Follmer, Alex Olwal, and Hiroshi Ishii. 2014. Physical telepresence: shape capture and display for embodied, computer-mediated remote collaboration. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 461–470.
28. William A McNeely. 1993. Robotic graphics: a new approach to force feedback for virtual reality. In *Virtual Reality Annual International Symposium, 1993., 1993 IEEE*. IEEE, 336–341.
29. Sumit Mehra, Peter Werkhoven, and Marcel Worring. 2006. Navigating on handheld displays: Dynamic versus static peephole navigation. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 4 (2006), 448–457.
30. K E Overvliet, J B J Smeets, and E Brenner. 2008. The use of proprioception and tactile information in haptic search. *Acta Psychologica* 129, 1 (Sept. 2008), 83–90.
31. Gian Pangaro, Dan Maynes-Aminzade, and Hiroshi Ishii. 2002. The actuated workbench: computer-controlled actuation in tabletop tangible interfaces. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*. ACM, 181–190.
32. Stephen Pheasant and Christine M Haslegrave. 2016. *Bodyspace: Anthropometry, ergonomics and the design of work*. CRC Press.
33. Myrthe A Plaisier, Astrid ML Kappers, Wouter M Bergmann Tiest, and Marc O Ernst. 2010. Visually Guided Haptic Search. *Haptics, IEEE Transactions on* 3, 1 (2010), 63–72.
34. Ivan Poupyrev, Tatsushi Nashida, Shigeaki Maruyama, Jun Rekimoto, and Yasufumi Yamaji. 2004. Lumen: interactive visual and shape display for calm computing. In *ACM SIGGRAPH 2004 Emerging technologies*. ACM, 17.
35. Roman Rädle, Hans-Christian Jetter, Simon Butscher, and Harald Reiterer. 2013. The effect of egocentric body movements on users' navigation performance and spatial memory in zoomable user interfaces. In *Proceedings of the 2013 ACM international conference on Interactive tabletops and surfaces*. ACM, 23–32.
36. R Rädle, H C Jetter, J Müller, and H Reiterer. 2014. Bigger is not always better: display size, performance, and task load during peephole map navigation. In *Proceedings of the 32nd . . .*
37. Majken K Rasmussen, Esben W Pedersen, Marianne G Petersen, and Kasper Hornbæk. 2012. *Shape-changing interfaces: a review of the design space and open research questions*. ACM.
38. Jan Richter, Bruce H Thomas, Maki Sugimoto, and Masahiko Inami. 2007. Remote active tangible interactions. In *Proceedings of the 1st international conference on Tangible and embedded interaction*. ACM, 39–42.
39. Eckard Riedenklaus, Thomas Hermann, and Helge Ritter. 2012. An integrated multi-modal actuated tangible user interface for distributed collaborative planning. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*. ACM, 169–174.
40. Michael Rohs, Johannes Schöning, Martin Raubal, Georg Essl, and Antonio Krüger. 2007. Map navigation with mobile devices: virtual versus physical movement with and without visual context. In *Proceedings of the 9th international conference on Multimodal interfaces*. ACM, 146–153.
41. Mark S Sanders and Ernest James McCormick. 1993. *Human factors in engineering and design*. McGraw-Hill book company.
42. Julian Seifert, Sebastian Boring, Christian Winkler, Florian Schaub, Fabian Schwab, Steffen Herrdum, Fabian Maier, Daniel Mayer, and Enrico Rukzio. 2014. Hover Pad: interacting with autonomous and self-actuated displays in space. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 139–147.
43. Makoto Shimojo, Masami Shinohara, and Yukio Fukui. 1999. Human shape recognition performance for 3D tactile display. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 29, 6 (1999), 637–644.
44. Mike Sinclair, Michel Pahud, and Hrvoje Benko. 2013. TouchMover: actuated 3D touchscreen with haptic feedback. In *Proceedings of the 2013 ACM international conference on Interactive tabletops and surfaces*. ACM, 287–296.
45. Jacqueline C Snow, Rafal M Skiba, Taylor L Coleman, and Marian E Berryhill. 2014. Real-world objects are more memorable than photographs of objects. *Frontiers in human neuroscience* 8 (2014).
46. Jürgen Steimle, Andreas Jordt, and Pattie Maes. 2013. Flexpad: highly flexible bending interactions for projected handheld displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 237–246.
47. Brygg Ullmer and Hiroshi Ishii. 1997. The metaDESK: models and prototypes for tangible user interfaces. In *Proceedings of the 10th annual ACM symposium on User interface software and technology*. ACM, 223–232.

48. Christopher R Wagner, Susan J Lederman, and Robert D Howe. 2002. A tactile shape display using RC servomotors. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2002. HAPTICS 2002. Proceedings. 10th Symposium on*. IEEE, 354–355.
49. Malte Weiss, Florian Schwarz, Simon Jakubowski, and Jan Borchers. 2010. Madgets: actuating widgets on interactive tabletops. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*. ACM, 293–302.
50. Yasuyoshi Yokokohji, Ralph L Hollis, and Takeo Kanade. 1996. What you can see is what you can feel-development of a visual/haptic interface to virtual environment. In *Virtual Reality Annual International Symposium, 1996., Proceedings of the IEEE 1996*. IEEE, 46–53.
51. Yiwei Zhao, Lawrence H Kim, Ye Wang, Mathieu Le Goc, and Sean Follmer. 2017. Robotic Assembly of Haptic Proxy Objects for Tangible Interaction and Virtual Reality. In *Proceedings of the 2017 ACM on Interactive Surfaces and Spaces*. ACM.