

# Tactile Line Drawings for Improved Shape Understanding in Blind and Visually Impaired Users

ATHINA PANOTOPOULOU, Dartmouth College and Boston University, USA

XIAOTING ZHANG and TAMMY QIU, Boston University, USA

XING-DONG YANG, Dartmouth College, USA

EMILY WHITING, Boston University, USA

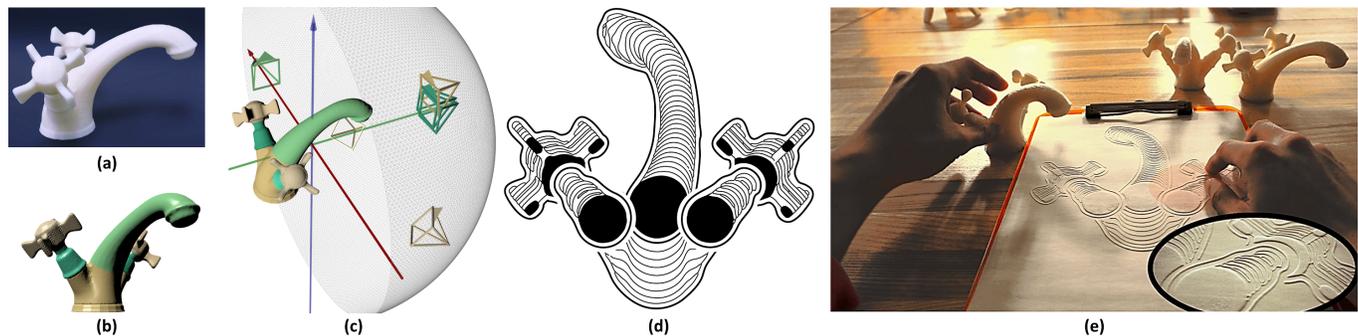


Fig. 1. We present a novel approach for generating tactile illustrations to improve shape understanding in blind individuals. (a) Physical 3D object (3D printed). (b) The input to our pipeline is a pre-partitioned object, colors indicate segmented parts. (c) A local camera is assigned to each part, with a master camera to combine the resulting multi-projection image. (d) Resulting stylized illustration. Cross-sections are used for texturing the interior of each part to communicate surface geometry. (e) We evaluated the technique in a user study with 20 blind participants. Tactile illustrations were fabricated using microcapsule paper.

Members of the blind and visually impaired community rely heavily on tactile illustrations – raised line graphics on paper that are felt by hand – to understand geometric ideas in school textbooks, depict a story in children’s books, or conceptualize exhibits in museums. However, these illustrations often fail to achieve their goals, in large part due to the lack of understanding in how 3D shapes can be represented in 2D projections. This paper describes a new technique to design tactile illustrations considering the needs of blind individuals. Successful illustration design of 3D objects presupposes identification and combination of important information in topology and geometry. We propose a twofold approach to improve shape understanding. First, we introduce a part-based multi-projection rendering strategy to display geometric information of 3D shapes, making use of canonical viewpoints and removing reliance on traditional perspective projections. Second, curvature information is extracted from cross sections and embedded as textures in our illustrations.

CCS Concepts: • **Human-centered computing** → *Accessibility systems and tools*; • **Computing methodologies** → *Shape analysis*; **Perception**; **Non-photorealistic rendering**.

Additional Key Words and Phrases: fabrication, design, accessibility, tactile shape perception, tactile images, non-photorealistic rendering

Authors’ addresses: Athina Panotopoulou, athina.panotopoulou@dartmouth.edu, Dartmouth College and Boston University, USA; Xiaoting Zhang, xiaoting@bu.edu; Tammy Qiu, tqiu@bu.edu; Emily Whiting, whiting@bu.edu, Boston University, USA; Xing-Dong Yang, xing-dong.yang@dartmouth.edu, Dartmouth College, USA.

© 2020 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3386569.3392388>.

## ACM Reference Format:

Athina Panotopoulou, Xiaoting Zhang, Tammy Qiu, Xing-Dong Yang, and Emily Whiting. 2020. Tactile Line Drawings for Improved Shape Understanding in Blind and Visually Impaired Users. *ACM Trans. Graph.* 39, 4, Article 89 (July 2020), 13 pages. <https://doi.org/10.1145/3386569.3392388>

## 1 INTRODUCTION

Understanding the 3D geometry of everyday objects via 2D media is essential to live, learn, and work. Designers communicate the shape of a product (e.g., a chair) via sketches and renderings. Consumers understand the design by browsing images of the product in a catalogue. Students learn physics, such as why a plane is shaped in a certain way for aerodynamics, from figures in a textbook. Visual perception of complex 3D geometry from a 2D projection is taken for granted by people with healthy vision. However, for people with near or total blindness, understanding 3D geometry of daily objects from current media is extremely challenging.

In this paper we introduce a novel approach to generating *tactile line drawings* that aid 3D shape understanding in users with near or total blindness, enabling the blind community to perceive complex 3D objects from 2D media. Members of the blind community rely heavily on tactile illustrations, defined as raised graphics (e.g. on paper) that are felt by hand. Common uses of tactile illustrations are to make visual information accessible in textbooks, maps [Brock et al. 2015], scientific diagrams [Brown and Hurst 2012], children’s literature [Claudet et al. 2008; Stangl et al. 2014; Stangl 2019], or museum exhibits. The tactile illustration in Figure 2 shows two representative samples, both used to give access to artifacts displayed in the Museum of Fine Arts Boston.



Fig. 2. Tactile illustrations (a,c) displayed alongside exhibits in the Museum of Fine Arts Boston (b,d) for accessibility purposes (Photographs ©Museum of Fine Arts, Boston, [www.mfa.org](http://www.mfa.org)). Black areas in illustrations indicate raised regions on the paper surface that can be perceived by touch.

To design these types of tactile illustrations, designers rely on guidelines provided by a number of associations, such as the Braille Authority of North America [2010]. Hundreds of pages of rules make manual design time consuming. For example, rules related to size, placement, variety, and form of different elements (e.g. lines, textures, points, labels, titles, and captions) permit easier exploration. Rules related to simplification, separation, elimination, and distortion of the illustration avoid clutter. Rules related to semantic segmentation, keying of the design, and selection of viewpoints permit easier understanding.

Unfortunately, studies have shown that shape understanding is flawed in traditional tactile diagrams, which can be seen from low object identification scores [Klatzky et al. 1993]. We aim to improve shape understanding by re-evaluating graphic design considerations and accounting for the unique needs of blind individuals. For example, while current tactile illustrations are created from a standard single viewpoint, a blind individual who has never had access to visual information might find it unnatural to explore a visual image that includes perspective distortions. This is supported by drawings created by blind individuals [Kennedy 1993], in which illustrations of 3D objects appear to use a “fold-out” method to arrange faces onto a flat surface [Kurze 1997].

In this paper we first present two formative studies that motivated our approach to tactile illustration. We investigate salient shape properties in tactile perception, and assess the strengths and weaknesses of existing tactile illustration styles. Next, we propose a novel method to render tactile graphics with a focus on improving shape understanding of 3D objects. We design a pipeline introducing a multi-projection rendering approach combined with texture infills that communicate surface geometry. Compared to a single-viewpoint illustration, our multi-projection method carefully places a local camera for each semantic part of the input object, so that each part may be rendered from its optimal viewpoint. A compositing stage aligns image layers created from each camera to maintain proper connectivity of the input object. Our texturing technique then uses closely spaced lines generated from cross sections to give cues about surface geometry of the object. In summary, we make the following contributions:

- We present results of a formative study involving sculpting 3D objects, aiming to discover salient information for blind users in haptic shape understanding (Sec. 3.1).

- We present results of a second formative study on the success of existing 2D tactile illustration guidelines for shape understanding (Sec. 3.2).
- We introduce a new design methodology for tactile illustrations based on the outcome of our formative studies, applying techniques in multi-projection rendering and geometry-aware textures (Sec. 4).
- We design and implement a user study evaluating the ability for blind users to understand 3D shape with our new illustration style (Sec. 5).

## 2 RELATED WORK

Our approach to tactile graphics is motivated by cognitive principles in tactile perception and design principles of stylized graphic illustration.

*Tactile Perception.* Mental representations in shape understanding are an active area of research in psychology and neurobiology, particularly concerning similarities in how humans perceive shapes visually and through touch [Farley Norman et al. 2004; James et al. 2002; Lakatos and Marks 1999]. A prominent theory on shape understanding is “recognition by components” [Biederman 1987; Blake and Sekuler 2006], suggesting that our visual system processes shape information by representing an object as a spatial arrangement of 3D parts called geons (generalized cones). Our visual system extracts geons and their arrangement from an image to recognize the object. Visual translation theory hypothesizes that haptic perception is translated to the same visual representations [Cooke et al. 2007; Erdogan et al. 2014; Handjaras et al. 2016; Lederman and Klatzky 2009], while other work suggests vision and haptics are functionally overlapping but do not necessarily have equivalent representations of 3D shapes [Farley Norman et al. 2004; James et al. 2002; Lakatos and Marks 1999]. Our proposed work on illustration style abides by visual shape recognition theory, where pictorial representations should clearly show subcomponents and their arrangement in space [Thompson et al. 2006].

*Tactile illustrations.* Tactile illustrations are typically made by expert designers in a time consuming process, involving identification and application of principles from existing guidelines (e.g. Braille Authority of North America [2010]). Rules relate to use of points, lines, textures, and labels, as well as clutter avoidance, e.g., by elimination, simplification, or separation of graphical elements. However, previous studies have still shown poor 3D shape perception from traditional 2D tactile diagrams. Similar to our formative study (Sec. 3.2), Klatzky et al. [1993] cite success rates at the level of 30% for identifying everyday objects, with some improvement observed for textured images compared to raised line graphics [Theurel et al. 2013]. Thompson and Chronicle [2006] introduced the TaxyForm system and discuss the need for re-evaluation of design methods.

Automatic tactile illustration generation has been studied for 2D input images, with improved edge extraction [Hernandez and Barner 2000], and filters for removing details too small to be perceived by touch [Way and Barner 1997]. In contrast, our proposed semi-automatic method operates on 3D geometric input and introduces a new stylization procedure. We were inspired by two

lines of work that explored non-standard illustration techniques: a) First, the procedure of “unfolding” or “flattening” 3D shapes to obtain 2D depictions [Kurze 1997], based on drawings made by blind individuals. We expand on this idea with a multi-projection strategy for general 3D objects to create flattened illustrations. Our approach considers viewpoint selection, known to play a key role in 3D shape recognition [Farley Norman et al. 2004]. b) Second, we expand on the use of line curvature and orientation, which have been shown to be salient features [Hsiao 2008; Yau et al. 2015], used successfully in 2D tactile Braille and Moon alphabets [Moore 2011]. Thompson et al. [2006] studied the use of curvature lines in shape perception, however, examples were constructed case by case and did not provide a system for application on novel objects. We develop a systematic approach to the design of tactile illustrations requiring minimal manual intervention.

More broadly, research on tactile graphics covers a range of applications such as maps and scientific diagrams [Brock et al. 2015; Brown and Hurst 2012], computer user interfaces [Li et al. 2019], children books [Cottin et al. 2008; Kim et al. 2015, 2014; Kim and Yeh 2015; Stangl et al. 2014; Stangl 2019; Walsh 2017], and artistic pieces [Ré 1981]. See Levesque [2005] for a survey.

*Stylized Rendering.* Extensive work has been done in computer graphics on non-photorealistic line drawings that improve how shape and topology are conveyed [Grabli et al. 2004]. E.g., using suggestive contours [DeCarlo et al. 2003], emulating pen strokes on parametric surfaces [Winkenbach and Salesin 1996], and silhouettes and hatching [Hertzmann and Zorin 2000; Kalogerakis et al. 2012; Praun et al. 2001; Zander et al. 2004]. Studies show the effectiveness of these techniques in depicting shape [Cole et al. 2009]. However these methods were designed for visual perception, and properties such as shading due to illumination are not relevant to tactile perception. Further, the cutaneous system for touch is limited in spatial resolution compared to the visual system [Blake and Sekuler 2006; Lederman and Klatzky 2009]. High frequency details are lost when perceiving haptically raised lines. Way and Barner [1997] identified size restrictions for the low resolution sense of touch which are enforced in most guidelines [Hasty 1999; McLennan et al. 1998; of North America 2010].

Geometric modifications can provide stylization [Liu and Jacobson 2019], or help visualize internal structure with approaches such as cutaways [Li et al. 2007], splitting objects [Islam et al. 2004], or using deformations [McGuffin et al. 2003]. Geometric modifications also help display the global shape of forms, by abstracting the geometry [Mehra et al. 2009], or improving visibility in exploded diagrams [Li et al. 2008; Tatzgern et al. 2010] for, e.g., complex mechanical assemblies. Non-geometric modifications include material changes, such as the use of transparency and ghosting [Diepstraten et al. 2003].

When representing 3D shapes in 2D media, choice of camera projections vary significantly depending on use case. For example, orthographic and oblique projections are used widely in CAD and technical drawing design [Ching and P. 2019; Krikke 2000], while linear convergent perspective projection is found more commonly, e.g., in videogames and film. Linear divergent perspective is used in artistic drawing [Howard and Allison 2011; JMS 2010], and non

linear perspective (e.g. curvilinear perspective [Sudarsanam et al. 2005]) projection characterizes wide angle photography and some artistic stylizations, e.g., Escher [Brosz et al. 2007]. Past work in multi-projection rendering [Agrawala et al. 2000; Yu and McMillan 2004] shows how images from different camera viewpoints can be combined into a single image. We draw from the multi-projection approach of Agrawala et al. [2000] and extend their work to using local cameras for segmented parts within the same object.

### 3 FORMATIVE STUDIES

In this section we present two formative studies that guided our new design approach for tactile illustrations. The first investigates what information is considered salient in tactile shape understanding. The second study assesses how successfully existing illustration guidelines convey 3D shape information.

#### 3.1 User Study 1: 3D Object Replication

We performed an object replication task to develop a deeper understanding of shape characteristics blind people consider salient. The experiment involved a sculpting task where the aim was to identify shape characteristics that participants focus on.



Fig. 3. Sculptures from User Study 1. (Top row) Reference objects. (Middle row) Participant 1 sculptures. (Bottom row) Participant 2 sculptures.

**3.1.1 Apparatus and Participants.** A variety of man-made and natural objects with different levels of shape complexity were selected. Two congenitally blind participants (1 female) participated in the study, where they were given an hour to sculpt five shapes, including a pyramid, dolphin, cup, chair, and human in a neutral standing pose (Fig. 3, top row) using 4 oz of polymer clay for each shape.

**3.1.2 Task and Procedure.** Participants were instructed to sculpt replicas of the shapes. They were also asked to identify each object and describe its shape and any difficulties encountered while sculpting.

**3.1.3 Observations.** The results of the study (see Fig. 3) suggest a number of interesting findings. First, the participants segmented the shapes into distinct geometric parts (F1). For example, they first created the four chair legs and brace segments, then proceeded to connect the pieces together. Participants were also able to correctly follow the topology of the reference shapes (F2). An example of

this is sculpting the correct number of pieces to construct the legs in the chair model. They were also able to sense simple shapes of the objects or the cross-section of a part (e.g., square base for the pyramid, rounded cross-section for the body of the dolphin, or cylindrical cross-section for the cup). Participants were also able to notice differences when cross-sections change within the same part (F3). For example, they noticed the change in cross-section from the front to the back of the dolphin body. Participants also pay attention to the curvature properties of the surface geometry (F4). For example, they flattened the facets of the pyramid and created distinct edges, or sculpted a rounded shape for the dolphin's body.

Participants encountered a number of challenges during the study. For example, participants were unable to accurately replicate positions of parts (F5). Sometimes they even misidentified relative positions (e.g., failing to make the dolphin's pectoral flippers symmetric). The participants also experienced difficulties in accurately capturing proportions (F6), e.g., the height-to-width ratio of the pyramid was less important than replicating a sharp point at the apex.

Our findings suggest that to communicate 3D shape information effectively to blind users, the following important characteristics should be conveyed through the illustration:

- Distinct parts within an object (F1)
- Topology of the shape (i.e. preserve part connectivity) (F2)
- Cross section shape and propagation (F3)
- Curvature of surface geometry (F4)

Our findings also suggest that the following characteristics have low importance in shape understanding:

- Accurate placement of the connection points (F5)
- Accurate proportions (F6)

Note that F5-F6 were independently observed, not inferred from F1-F4. The preliminary findings in our object replication study align well with existing theory on shape understanding, in particular the "recognition by components" theory [Biederman 1987; Blake and Sekuler 2006] of how any view of an object can be represented as a spatial arrangement of 3D parts (called geons, or generalized cones).

## 3.2 User Study 2: Assessment of Existing Illustration Guidelines

**3.2.1 Apparatus and Participants.** Seven congenitally blind participants (3 female, aged between 30 and 72) participated in our study. We chose five reference objects familiar to our participants: teapot, lamp, eyeglasses, chair, and rocking horse. We created tactile illustrations for each object in three different styles (Fig. 4):

- (1) Line drawing with object in non-canonical view (also known as perspective or 3/4 view).
- (2) Line drawing with object in canonical view (e.g. front or side).
- (3) Line drawing with canonical view and texture-infills to indicate different semantic regions in the illustration.

In case (2) we used two viewpoints (i.e. 2 corresponding illustrations) if important information was contained in both, following recommendations by the Braille Authority of North America (BANA) guidelines [2010]. In case (3) the separation into semantic regions and selection of textures were following BANA guidelines.

All illustrations were made by a trained graphic designer. The physical tactile diagrams were created by laser engraving commercially available capsule paper.

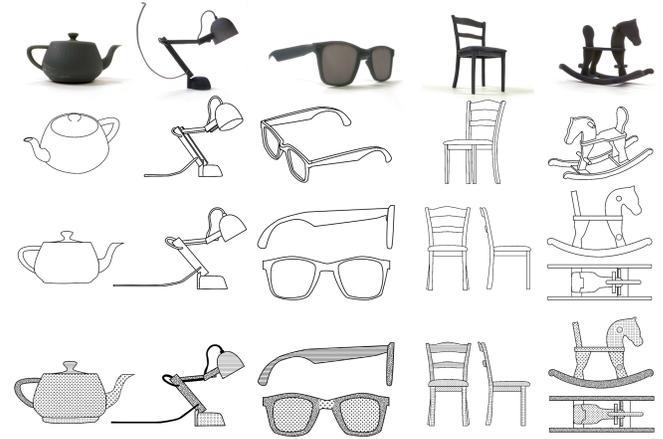


Fig. 4. 3D printed reference objects for User Study 2 (row 1). Corresponding illustrations: non-canonical views (row 2), canonical views (row 3), and textured canonical views (row 4).

**3.2.2 Task and Procedure.** During the study, participants were asked to use the three tactile illustrations to recognize the tested objects, identify parts of the objects (e.g., locate the spout of the teapot), explain the perceived shape of the parts, explain the connection between the parts (e.g., where the spout connects to the main vessel), and count the number of the parts. We showed the objects in the same sequence to all users: (1) non canonical, (2) canonical, (3) textured canonical.

**3.2.3 Observations.** Several important findings were observed from our study regarding the key characteristics related to tactile line-drawings in communicating 3D shape information, including the number of viewpoints, viewpoint selection, occlusion, and shape understanding.

**Multiple viewpoints.** According to the Braille Authority of North America (BANA) guidelines for tactile graphics [2010], multiple views should be provided if they contain information important for the intended task (guidelines: 10.1.4; 7.1.1.5; 3.6.2). However, we found people could not successfully identify correspondences between the different illustrations (F7). Instead, participants often chose to explore their preferred viewpoint to perform a task, even if the essential information was contained in the other view. It was considered cumbersome to explore more than one diagram at a time.

**Viewpoint Selection.** Selection of viewpoints has been shown to influence object recognition accuracy [Heller et al. 2009]. We also found that participants' ability to identify an object from an illustration had higher success rates for canonical viewpoints (21%) vs. non-canonical viewpoints (6%) (F8). However, participants were generally not satisfied with any of the canonical viewpoints provided. For example, in the eyeglasses illustrations participants commented that the front view depicted the correct size for the lenses but the

wrong size for the frame arms, and the opposite was true for the side view. A participant suggested combining viewpoints: showing the lenses from the front and the frame arms from the side. That is, create a combined multi-projection image in which each part will be shown in its optimal view (F9).

*Occlusions.* One of the most common complaints from participants was the absence of parts from the diagrams, which was a consequence of occlusion (F10). For example, in the chair illustrations, front and side views were provided (see Fig. 4) and a participant indicated that the chair seemed to have 2 legs instead of 4. This is consistent with an existing study by Moringen et al. [2017] that showed occlusions affect tactile shape recognition.

*Shape understanding.* Illustrations rendered in canonical views were generally better for understanding surface geometry compared to non-canonical views (F8). When asked about parts being more rectangular vs. cylindrical, users were 71% correct for canonical illustrations vs. 48% for non-canonical views. However, in some cases the non-canonical views improved shape understanding. For example, the non-canonical view of the teapot was more successful for recognizing the spherical shape of the main container, whereas in the canonical view it was perceived as cylindrical or cubical. Possible reasons are the curved edge along the bottom of the teapot and ellipsoidal shape of the lid (see Fig. 4). This is in agreement with a study by Thompson et al. [2006] which showed that curved lines as textures in tactile illustrations give indication about the geometry of the represented shape. Further, a study by Yau et al. [2015] showed that line orientations and the level of curvature and curvature direction of lines changes are identifiable by touch (F11).

**3.2.4 Foundations for the new design.** Based on User Study 2, we propose a number of additional design guidelines to facilitate 3D shape understanding using tactile illustrations:

- Contain all information in a single illustration (F7)
- Depict canonical viewpoints (F8)
- Combine multiple viewpoints within a single illustration so that each part of the object can be depicted optimally (F9)
- Choose a viewpoint of the object to reduce occlusions (F10)
- Use curved lines to communicate surface geometry (F11)

## 4 TACTILE ILLUSTRATION METHOD

### 4.1 Overview

Given the design foundations derived from our formative studies, we present a new approach to creating tactile illustrations of 3D objects. Our method combines images from different viewpoints, applying techniques in multi-projection rendering. We also introduce a new approach to generating tactile infill textures that communicate the surface geometry of the object.

Our illustration approach draws inspiration from the artistic multi-projection rendering stylization introduced by Agrawala et al. [2000]. Similar to their work, we create a set of local cameras which render images from different viewpoints. We then employ a master camera to combine the local renderings into a single layered image. A key difference to Agrawala et al. [2000] is that our local cameras are assigned to segmented parts within the same object.

This means our method must constrain relative positions of part-level renderings in order to maintain correct connectivity in the final illustration. Further, our approach automates camera positioning by defining a notion of optimal viewpoint for each part.

The input to our pipeline is a partitioned 3D object. Users can either manually partition the object or use an automated technique (e.g. [Zhou et al. 2015]). We use the PartNet dataset [Mo et al. 2019] consisting of pre-partitioned 3D objects in a wide range of common object categories, and provided in up-right orientation. Our pipeline outputs a stylized multi-projection rendering, which can then be fabricated to produce a tactile illustration. The main steps of our algorithm are as follows:

- (1) Determine local camera poses. A *part camera* is assigned to generate a rendering for each decomposed part of the object, with position and orientation chosen according to visibility objectives (Sec. 4.3). A *master camera* is employed to combine the multiple renderings into a unified illustration (Sec. 4.2).
- (2) Generate texture infills for each part rendering. We integrate surface geometry information using cross-sections extracted along the object skeleton (Sec. 4.4).
- (3) Compositing the final illustration from local part renderings. The combination of image layers accounts for relative up-vector alignment (Sec. 4.5), part connectivity within the 3D object, and occlusion handling (Sec. 4.6).

### 4.2 Master camera placement

In our multi-projection approach a master camera is used to organize the rendered images from each local part camera. As discussed in Sec. 3.2.4 (design foundation (F10)), viewpoints that result in non coincidence of parts on the image (i.e., not more than one part appearing on each pixel) are preferred. The view from the master camera helps to define the local part camera frames by prioritizing visibility in the final composition.

To determine the master view direction  $\hat{\mathbf{d}}_M$  we follow the approach of Secord et al. [2011], where several attributes are defined for identifying preferred visual viewpoints. We notice that viewpoints that maximize the silhouette length result in decreased part coincidence. Silhouette length is defined as the overall length of the object's silhouette in the camera's image plane. E.g., in Fig. 6(D), overlap between the handle and main container is small when viewed from a direction that maximizes silhouette length.

Next, we align the up-vector of the camera  $\hat{\mathbf{u}}_M$  with the up-vector of the object. This results in a top to bottom layout in the final illustration (i.e., the base of the object is positioned at the bottom of the illustration), which is a natural choice and follows the Braille Authority of North America (BANA) guidelines and standards for tactile graphics [2010] (6.11). Note the object up-vector is not considered in Secord et al. [2011] for viewpoint selection. We fix the up-vector and constrain the camera location to a ring centered in the object (see Fig. 5(b)).

### 4.3 Part Camera Placement

Next, we determine a local frame for each part camera. As discussed in Sections 3.2 and 3.2.4 (design foundation (F8)), canonical views tend to improve shape understanding for blind users. To identify

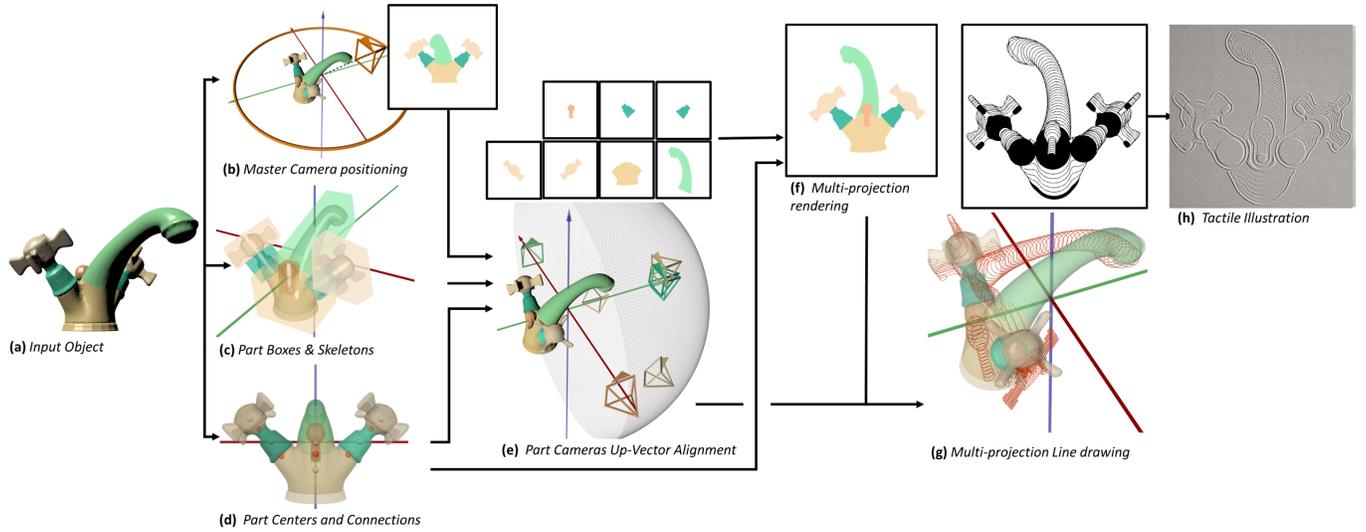


Fig. 5. Overview of our pipeline for the design of tactile illustrations. (a) Input model, colors represent different parts of the faucet. (b) Master camera placement (inset box shows camera view). (e) Local cameras are assigned to each part (inset boxes show camera views). (f) The image layers are composited; positioning of each layer is guided by the master camera and the part connectivities (d). (g) Infill textures are generated based on cross sections of each part (black represents raised areas). (h) The fabricated illustration created with microcapsule paper.

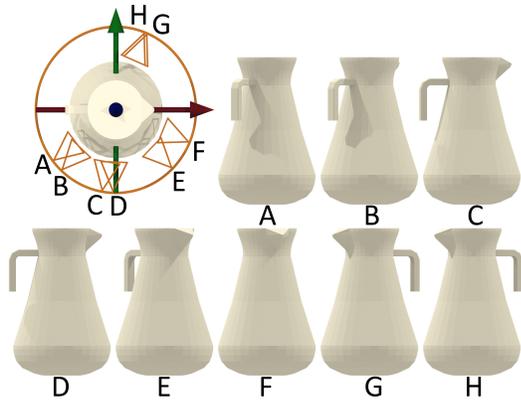


Fig. 6. Different metrics considered for finding the master camera placement. The viewpoints maximize: (A) Silhouette Curvature Extrema, (B) Silhouette Curvature, (C) Projected Area, (D) Silhouette Length, (E) Surface Visibility, (F) Maximum Depth, (G) Depth Distribution, (H) Viewpoint Entropy. The maximum silhouette length (D) decreases part occlusion.

the canonical viewpoint for each part we first extract a coordinate frame  $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3\}$  using PCA, similar to Vranic et al. [2001].

We next extract a one dimensional skeleton for each part using the mean curvature skeleton by Tagliasacchi et al. [2012] (see white lines in Fig. 5(c)). We reject the PCA axis that most closely aligns with the skeleton, which will be important later when creating infill textures.

Between the remaining axes we choose the direction corresponding to the largest projected area. For simplicity this is done heuristically by aligning a bounding box to the frame (see Fig. 5(c)), and selecting the axis corresponding to the largest face on the bounding

box (i.e. face with normal  $\hat{\mathbf{e}}_i$ ). Between the  $\pm\hat{\mathbf{e}}_i$  directions, we choose the one that gives a viewing direction in the same half space as the master camera, where local camera viewing direction is  $\hat{\mathbf{d}}_P = \mp\hat{\mathbf{e}}_i$ . The part camera location is given by  $\mathbf{p}_P = \mathbf{c}_P - r\hat{\mathbf{d}}_P$ , where  $\mathbf{c}_P \in \mathbb{R}^3$  is the centroid of part  $P$  and  $r$  is the distance to the camera ( $r > 0$ ). The part  $P$  is centered in the local camera viewport.

#### 4.4 Cross Section Textures

The next step is to create the infill textures. Compared to the BANA standard [2010] which uses uniform textures to denote semantic regions, we choose to incorporate properties of surface geometry motivated by design foundations (F3, F4, F11). Many methods have been explored for creating interior lines and textures in non-photorealistic rendering, such as hatching and cross-hatching [Kalogerakis et al. 2012] or apparent ridges [Judd et al. 2007]. However these approaches are based on illumination effects of shading or view-dependent properties which are not relevant to blind users for tactile perception of 3D forms. Similar to Deussen et al. [1999] we choose to create texture infills using cross section boundaries to illustrate the geometry of each part.

As noted in Sec. 4.3, for each part we extract a 1D skeleton using the method of Tagliasacchi et al. [2012]. We initialize the cross sections to be uniformly sampled along the skeleton, with normals,  $\hat{\mathbf{n}}_{CS}$ , aligned with the local skeleton direction. Next, to allow for visibility of cross-sections on the image and remove perspective distortion, each cross section is flattened by rotating to lie parallel to its corresponding part camera view plane (see Fig. 5(g) and Fig. 7). The skeleton sample point remains fixed with the rotation axis given by  $\hat{\mathbf{d}}_P \times \hat{\mathbf{n}}_{CS}$ , where  $\hat{\mathbf{d}}_P$  is the view plane normal for local part camera  $P$ . We then render a line for each cross section silhouette from the corresponding part camera. The layering of cross section silhouettes

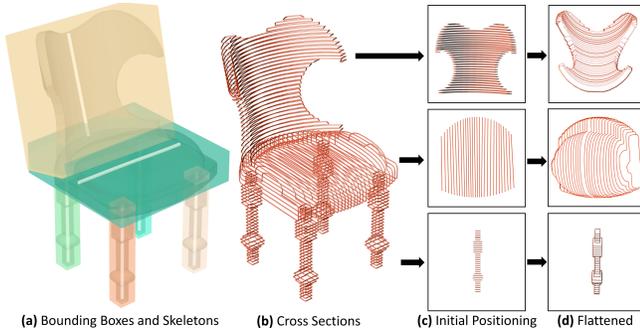


Fig. 7. The chair displaying the bounding box and skeleton for each part (a). Cross sections sampled uniformly along the skeleton. The cross sections of each part as seen from their corresponding local cameras (c), and after the flattening operation (d).

in the image follows the sampling ordering along the skeleton (see pseudocode in Alg. 1).

The spacing of cross section lines considers limits on resolution for tactile perception. The cutaneous system has poor spatial resolution compared to vision [Blake and Sekuler 2006; Lederman and Klatzky 2009; Way and Barner 1997], such that high frequency details are lost when haptically perceiving raised lines. We follow recommendations from the BANA guidelines [2010] of 2-3mm gaps between elements to be distinguishable by hand.

---

#### ALGORITHM 1: FLATTENEDCROSSSECTIONS ( $Sk, Shape_P, \hat{d}_P$ )

---

**Input:**  $Sk = [sk_1 \dots sk_n]$  a sequence of regularly sampled points on the skeleton. The geometry of the part,  $Shape_P$ . View plane normal  $\hat{d}_P$  for camera of part  $P$ .

**Output:** Cross sections  $CS = [cs_1 \dots cs_n]$  parallel to view plane and layered. Each  $cs$  defined by sample point from skeleton  $b_{cs}$ , normal  $\hat{n}_{cs}$ , and  $Boundary_{cs}$ .

```

1 for  $sk_i \in Sk$  do // Extract cross sections along skeleton
2    $b_{cs_i} \leftarrow sk_i$ 
3    $\hat{n}_{cs_i} \leftarrow$  local tangent direction at  $sk_i$ 
4    $Boundary_{cs_i} \leftarrow Shape_P \cap CUTTINGPLANE(b_{cs_i}, \hat{n}_{cs_i})$ 
5    $cs_i \leftarrow CROSSSECTION(b_{cs_i}, \hat{n}_{cs_i}, Boundary_{cs_i})$ 
6 end
7 for  $cs_i \in CS$  do // Align cross sections with view plane
8   if  $\hat{d}_P \neq \hat{n}_{cs_i}$  then
9      $RotAxis \leftarrow \hat{d}_P \times \hat{n}_{cs_i}$ 
10     $PivotPt \leftarrow b_{cs_i}$ 
11     $RotAngle \leftarrow ANGLE(-\hat{d}_P, \hat{n}_{cs_i})$ 
12     $cs_i \leftarrow ROTATECROSSSEC(cs_i, RotAxis, PivotPt, RotAngle)$ 
13  end
14   $b_{cs_i} \leftarrow b_{cs_i} - SmallStep * i * \hat{d}_P$ 
15 end
16 return CS
```

---

#### 4.5 Part Camera Up-Vector Alignment

We next align the orientation of the local part cameras. In Section 4.3 we determined the view direction for each part camera,  $\hat{d}_P$ , prioritizing the direction of max projected area for the part. However, the up-vector orientation of the camera,  $\hat{u}_P$ , was left ambiguous. Our approach takes into account the master camera viewpoint and

connectivity of parts within the 3D object. As an initialization we use the vector field provided by Lopes et al. [2013] – the vertical axis is the master camera up-vector,  $\hat{u}_M$ , the normal vector corresponds to the local camera view direction,  $\hat{d}_P$ , and the tangent given by the vector field provides an initial guess for  $\hat{u}_P$ .

The next step is to identify connection points between adjacent parts. For each part  $P_j$ , we collect the set of adjacent parts  $\mathcal{A}_j$ . We assign a representative point  $\mathbf{a}_{ij} \in \mathbb{R}^3$  to each connecting part  $P_i$ , for  $i \in \mathcal{A}_j$ , computed as the centroid of the contact surface between  $P_j$  and  $P_i$ . The part centroids  $\mathbf{c}_j$  and connection points  $\mathbf{a}_{ij}$  are shown for the faucet in Figure 5(d).

We make a slight adjustment for the reference “center” point of the part,  $\bar{\mathbf{c}}_j$ . If there is only one connection point, the center of the part  $\bar{\mathbf{c}}_j$  coincides with the mesh centroid. Otherwise the average of all connection points is used. We then compute a rotation for each part camera’s up-vector,  $\hat{u}_j$ , such that the relative *horizontal* positioning between  $\bar{\mathbf{c}}_j$  and  $\mathbf{a}_{ij}$  is the same when viewed in the master and part camera image planes (see Fig. 5(f), 8). To achieve this we use the following steps:

Recall that we are using orthographic cameras, where the object size in the image plane does not diminish with depth. We define the vector  $\mathbf{v} = \mathbf{a}_{ij} - \bar{\mathbf{c}}_j$  in world space. The horizontal distance in the master camera image is given by  $w = \mathbf{v} \cdot \hat{\mathbf{x}}_M$ , where  $\hat{\mathbf{x}}_M$  is the  $x$ -axis of the master camera viewport. We then determine the rotation angle  $\theta_{ij}$  of the camera frame around the view direction  $\hat{d}_j$ . The angle is chosen such that  $\mathbf{v} \cdot \hat{\mathbf{x}}'_j = w$ , where  $\hat{\mathbf{x}}'_j$  is the  $x$ -axis of the viewport for part camera  $j$  after rotation (see Fig. 8).

We repeat the same process for all connection points of the part and use the average rotation,  $\bar{\theta}_j = avg(\theta_{ij})_{i \in \mathcal{A}_j}$  as the final alignment of the up vector,  $\hat{u}'_j = \mathbf{R}(-\bar{\theta}_j)\hat{u}_j$ . The rotation axis is the viewing direction  $\hat{d}_j$ .

#### 4.6 Compositing Part Renderings

The final stage is to combine the images from each local part camera to create the composited illustration. We compute a depth value for each part, calculated as the distance from the master camera location,  $\mathbf{p}_M$ , to the furthest vertex of the part’s bounding box (in the master camera coordinate frame). The order controls the layering

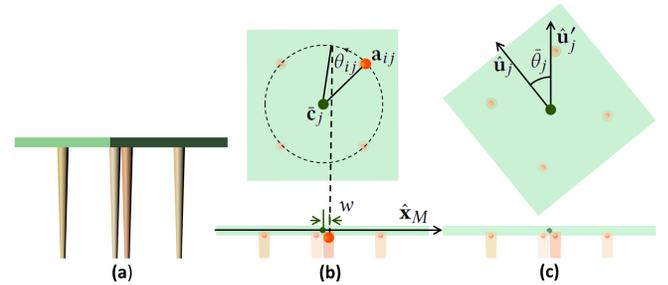


Fig. 8. Alignment of the part camera up-vector. (a) Input object as seen from the master camera. (b) Center and connections for the tabletop from the part camera view (top) and master camera view (bottom). Angle computed so that horizontal distances of connections,  $w$ , match the master camera view. (c) Final upright orientation of the local camera,  $\hat{u}'_j$ .

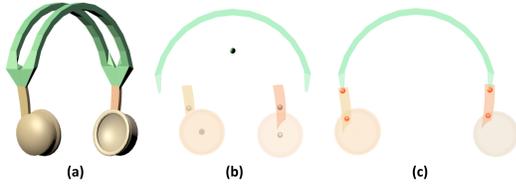


Fig. 9. Aligning part camera image layers. (a) The input model as seen from the master camera. (b) If the centroids in the master camera view are used for aligning part renderings, the resulting image is disconnected. (c) Instead, the connection points are used for aligning image layers.

with which images appear on the multi-projection illustration. The default is to start from the farthest and move to the closest part. The user may choose to change the ordering manually since this does not handle many occlusion cases, e.g., where parts close to the camera completely occlude other parts.

Initially all image layers are stacked such that their local frames align (i.e., image centers coincide). We traverse the graph of all parts starting with the farthest part, where nodes are the parts  $P_j$ , and edges are their part connectivity  $\mathcal{A}_{ij}$ . When we traverse the edge  $ij$ , i.e., a connection from part  $P_j$  to its adjacent part  $P_i$ , we update the position of the image layer from part camera  $i$ . The image layer is translated to enforce the part renderings to align at their connection point (Fig. 9). The translation vector is given by  $\mathbf{t}_j \in \mathbb{R}^2$ :

$$\mathbf{t}_j = \text{proj}_j(\mathbf{a}_{ij}) - \text{proj}_i(\mathbf{a}_{ij})$$

where  $\text{proj}_i(\mathbf{a}_{ij})$  is the projection of connection point  $\mathbf{a}_{ij}$  in the image plane of part camera  $i$ . After applying the transformations, the connection points from different camera renderings align on the image plane. Note we assume no cycles exist, this procedure does not handle cases where the connectivity of parts forms a multiply-connected graph rather than a tree.

As a final post-processing step, we create an outline around each of the part textures of 1mm thickness, and enforce a 3mm wide gap between each part (see implementation details in Sec. 6). This helps to differentiate between parts, following BANA guidelines [2010]. Renderings for the tactile illustrations of a selection of objects can be seen in Fig. 5(g) and Fig. 10.

## 5 USER STUDY 3: EVALUATION

We conducted a user study to evaluate the effectiveness of our novel tactile diagrams. The study included an object recognition task, in which participants with near or total blindness were invited to a controlled lab environment to identify a set of daily objects through the proposed tactile illustrations.

**5.0.1 Apparatus and Participants.** Our study consisted of 20 blind participants: 12 female, aged between 26 and 72, and 14 congenitally blind. Seven objects were used in the study, we chose common household objects that would be familiar to the participants (in the order presented): monitor, square table, headphones, chair, curved monitor, round table, and faucet (See Fig. 10). For each object, the study materials included:

- (1) Three variations of the 3D reference object (Fig. 12). All three versions were physically fabricated using 3D printing.

- (2) One tactile illustration generated with our technique (Fig. 10). Fabricated with microcapsule paper (Fig. 11).
- (3) One tactile illustration designed following BANA guidelines as a baseline. Fabricated with microcapsule paper.

Including variations of the 3D objects allowed us to measure how well our illustration technique helps participants precisely perceive the 3D geometry information and recognize subtle differences between the reference objects and their variations. To create variations of the 3D reference objects, each model was modified by either a replacement operation, subtraction, or addition of parts. The variations of the objects are as follows (see Fig. 12):

- Monitor 1: changes in base design
- Table 1: leg design (rectangular vs. cylindrical cross section, and cylindrical with connectors)
- Headphones: size of ear cushion and outer shape of earpieces
- Chair: curved vs. flat back, and presence of cushion
- Monitor 2: curved vs. flat screen and CRT style
- Table 2: number of legs and straight vs. tilted shape
- Faucet: number and shape of lever arms on handles
- Jug (training example): cross section shape of main container

Example materials for the chair are shown in Figure 13. Materials for all objects in the study are included in Supplemental Materials.

**5.0.2 Task and Procedure.** The tactile illustrations were presented to our participants, whose task was to find the corresponding 3D reference object. A baseline was included to allow us to measure the performance of our technique over the state-of-the-art. The baseline was produced by following the Braille Authority of North America (BANA) guidelines [2010]. See sample materials for the chair object in Fig. 13. Participants were trained in how to interpret the tactile illustrations using the jug example. The procedure was as follows:

- (1) Participants were given materials for one object at a time consisting of three variations of the object (3D printed), and one illustration of “type A”.
- (2) The participant was asked to choose which one of the three object variations most closely matched the illustration.
- (3) Steps (1) and (2) were repeated for all seven objects.
- (4) Steps (1-3) were then repeated using illustrations of “type B”.

For half of the participants “type A” was the baseline (BANA guidelines) and “type B” was our illustrations. For the other half of participants the illustration type was reversed. We chose to use only three variations of the reference object to avoid cognitive overload, this was determined after running a small pilot study.

Our quantitative measures included matching rate, which measures the number of times participants were able to correctly match a tactile illustration with the corresponding 3D object. Qualitative observations were also collected to assess how easy it was for participants to understand the 3D shape information shown in the tactile illustrations and identify the subtle differences between them.

**5.0.3 Observations.** We analyzed the data using a generalized linear mixed model for a comparison between the baseline BANA illustrations and our technique, considering each response an individual sample. The network includes two fixed effects: design and object, and a random effect: participant, and is built on a family of binomial distributions. The variable object has seven levels (monitor 1, table



Fig. 10. Resulting designs for tactile illustrations using our technique, with 3D input models shown.

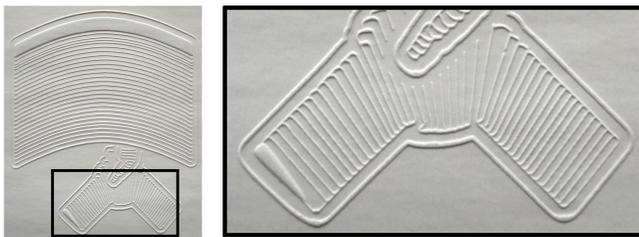


Fig. 11. Photo of tactile illustration of the monitor, including close-up view of the base. Fabricated with microcapsule paper.

1, headphones, chair, monitor 2, table 2, and faucet) and the design has two (BANA, our technique). To compare success rates within each object we do Post-Hoc pairwise comparisons on estimated marginal means and apply Bonferroni corrections [Lenth 2019].

The mean success rate in matching the correct 3D reference object to the tactile illustration was 58% for our technique versus 29% for the BANA baseline. We found that there exists a significant difference in shape understanding among the two types of illustrations ( $\chi^2(1, N = 140) = 22.77, p < 0.005$ ). Figure 14 shows the mean success rate over all objects, as well as per object success rates. Confidence intervals are computed by standard error. There is also a significant effect of Object ( $\chi^2(6, N = 140) = 31.55, p < 0.005$ ). Post-hoc analysis shows that higher success rates were observed for our technique for the majority of objects, with monitor 1, chair, and faucet remarkably better, with an improvement of  $\geq 45\%$  in mean success rate (all  $p < 0.005$ ). Table 1, monitor 2, and table 2 also have large improvement (e.g., improvement of  $\geq 15\%$  in mean success rate), but without statistical significance. An exception to this trend is the headphones where participants responded more accurately using the BANA baseline ( $p < 0.05$ ). Figure 15 compares results between congenitally (14) and late (6) blind participants. Both groups show better success rate with our technique.

Although our results are promising, there were still misconceptions caused by our technique. One observation was that overlapping parts were sometimes overlooked. For example, in our illustration for table 1 (square top, see Fig. 10) one of the legs is entirely

contained by the tabletop. Several participants were unable to locate this leg, likely since the silhouette was not easily discoverable.

The headphones were the one object with poorer mean success rate than the baseline. We found that participants were not using the infill textures to understand the shape of the earpieces, rather they were judging shape just by the outline which may have indicated a spherical instead of cylindrical shape. Another consideration was the lack of symmetry in our illustration – although identical, the two earpieces were shown from different views which produced different silhouettes and seemed to confuse participants.

In order to understand user perception of our new illustration style, we conducted a usability survey based on the System Usability Scale (SUS) [Brooke 1996]. With a maximum score of 100 we found similar scores on average of 65 for our technique and 62 for the BANA baseline. This result also indicates that participants tended to be open to new ideas and solutions for the improvement of tactile illustrations.

At the end of the study we collected additional unstructured feedback from participants, summarized here:

*Infill textures.* Participants had positive feedback on the infill textures for both our technique and the BANA baseline approach. Participants found the BANA textures useful for differentiating between different parts (2 users). In contrast, our textures were useful for understanding shape details (2 users).

Participants generally liked the textures in our illustration technique, with two users saying it is a consistent way to represent shapes, and four users saying the shape could be felt clearer/easier. Other positive remarks were that the designs were “more detailed,” “interesting,” and that the textures were “more tactile.” Specific use cases were suggested, including art objects in a museum (e.g., statues), and for depicting architecture (e.g., Cape vs. Victorian roof styles or different pillar styles). Three users raised concerns, mostly related to the clutter that the lines created, which made the illustration feel “complicated” and “everywhere the same,” so that they could not easily identify “where one piece starts and where it stops.”

*Multiple viewpoints.* Participants generally appreciated that our technique created a single unified illustration, “there was only one shape to look at.” Two participants commented that our technique

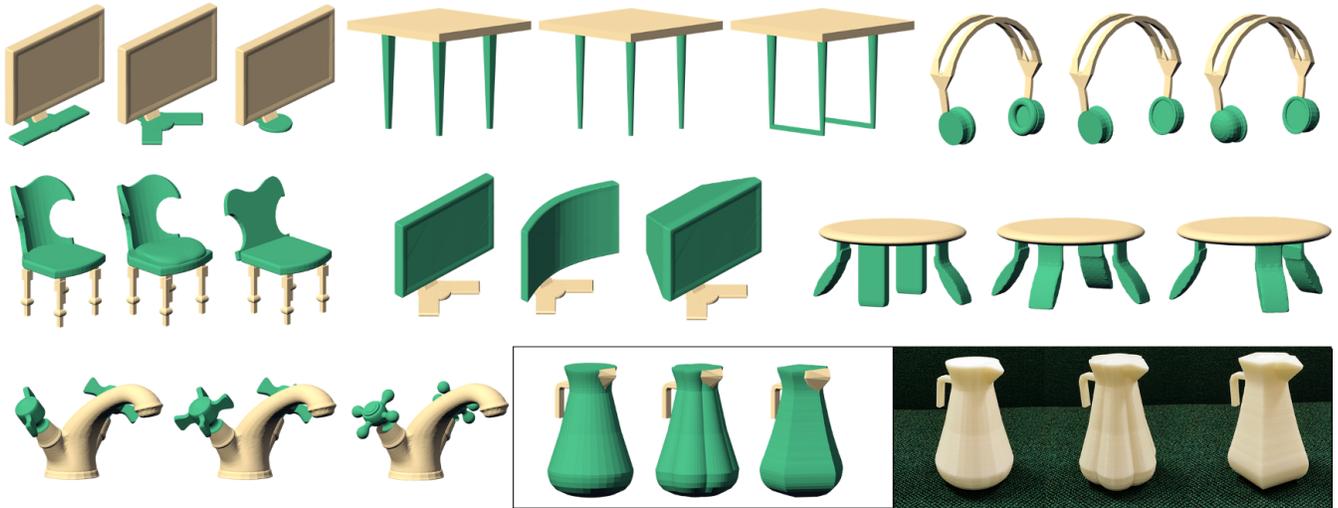


Fig. 12. For each 3D object included in User Study 3: Evaluation, we created three versions with minor geometric variations. Beige regions are identical among the three versions; green regions have been modified. The models were 3D printed so that participants could compare (by touch) the 3D objects depicted in the 2D tactile illustrations. The tactile illustrations shown in Fig. 10 correspond to the middle version in each case.

represents the shape more effectively, is more simplified, and conveys the object more efficiently. One participant commented: “If I had seen the diagrams [our technique] and not the objects, but I was told that the diagram was a chair, then I could have interpreted it. But once we moved to different angled views [BANA baseline] that became less and less something I would be confident about.” However, three users found difficulties in interpreting our illustrations, saying they looked more like an “artist drawing” and “more 3-dimensional,” and not knowing the view was problematic.

Regarding the BANA baseline illustrations, four participants commented that it was difficult to explore multiple views (e.g., top and side) which they later needed to match and compare. E.g.: “in some of the views you did not see all the object, you did not understand what you were looking at, you had to use both of them and kind of put them together or use what information you could get from the one that made the most sense.” This response highlights the benefits of a single illustration.

*Fabrication issues.* One participant noted that he would like not only thickness differences in lines, but also height differences, while another commented that he would prefer sharper lines.

## 6 FABRICATION AND IMPLEMENTATION DETAILS

*Implementation Details.* The skeleton extraction uses the CGAL implementation of the mean curvature skeleton [Tagliasacchi et al. 2012]. To create the renderings for User Study 2, we extracted lines using Freestyle [Grabli et al. 2008], making modifications manually when needed in Adobe Illustrator. We created lines of 1mm thickness with 3mm wide gaps between different object parts and then modified according to BANA guidelines [2010]. To create the BANA baseline renderings for User Study 3, we follow the same approach, using five pre-made textures tested with a blind individual for their identifiability. When creating illustrations with our technique for

User Study 3, we created an outline around each part texture using Adobe Illustrator. We set the texture line thickness at 0.5mm, and the outlines and gaps are the same as User Study 2. As some manual steps were used in creating the outlines, this post-processing stage prevented large-scale evaluation of our pipeline.

*3D Reference Objects.* The 3D objects in User Study 2 were selected from online repositories: archibaseplanet.com, archive3d.net, and www.thingiverse.com. For User Study 3, the 3D reference objects were selected from the PartNet dataset [Mo et al. 2019]. PartNet objects are pre-partitioned and provided in up-right orientation. Variations of the objects were created using Blender 2.79.

*Fabrication Details.* The objects in User Study 2 were fabricated with a Form 3 SLA 3D printer. The objects in User Study 3 were fabricated with a UPrint FDM 3D printer and Form 3 SLA 3D printer. The same printer was used for all variations of an object for consistency.

To fabricate the tactile illustrations we used microcapsule paper. The capsules expand under heat application, which was applied with laser engraving. We used an Epilog laser cutter with settings: 60 Watts, 100 speed, and power ranging between 12 to 15. Alternatively: 40 Watts and speed ranging between 16 to 20. Each image takes 20 to 30 minutes to complete.

## 7 DISCUSSION

### 7.1 Limitations & Future Work

We have presented a new methodology for generating tactile illustrations. We introduced a preliminary algorithm that integrates design considerations on tactile graphics motivated by a set of formative studies. However, many opportunities exist for improving the technique and experimenting with alternative stylizations. For example, further investigation could be done into selecting optimal viewpoints for parts, and layering images appropriately to handle



Fig. 13. Example materials for User Study 3: Evaluation. Participants are given three variations of the object (row 1), produced by 3D printing (row 2). They are given one tactile illustration style – either our new multi-projection style (row 3, left) or the BANA baseline style (row 3, middle-right). Illustrations are produced using microcapsule paper (row 4). Participants are asked to identify which of the three variations matches the illustration. (See Supplemental Materials for all materials used in study 3).

occlusions. We also hope to scale up to more complex objects while maintaining a usable level of detail. Future work should investigate how shapes that are not approximated well by 1D skeletons can be represented (e.g. objects with concavities such as bowls). Future work can also develop treatment for objects with cycles. Currently, each part is positioned to coincide with its parent connection, and a part that closes a cycle is not taken into account.

Furthermore, symmetry has been shown to be a key characteristic of tactile shape perception [Bauer et al. 2015], as our user study also indicates, and should be integrated into the design process in future work. Similarly, our method could incorporate special handling of repetitions within an object to communicate identical parts. Additionally, our method could be developed to incorporate other tactile

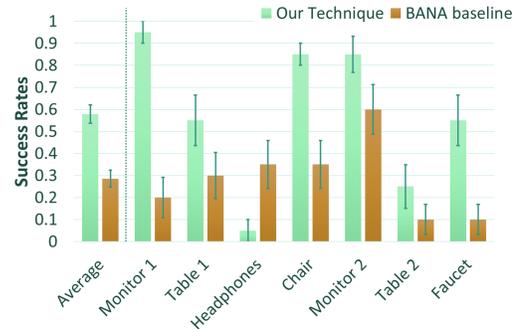


Fig. 14. User Study 3 results, shows mean success rate for matching tactile illustration to correct 3D object. Results shown for average over all objects, and per object success rates.

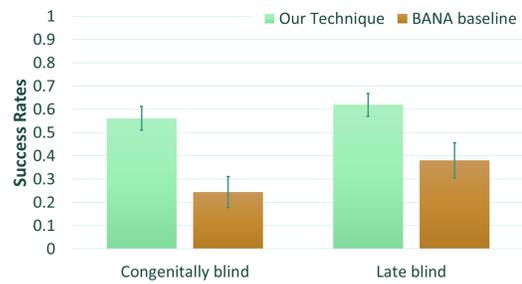


Fig. 15. User Study 3 results, shows mean success rate for matching tactile illustration to correct 3D object. Grouped by age of vision loss (self-identified).

saliency properties, such as the functional tactile saliency metrics developed by Lau et al. [2016] which was based on grasping.

Our illustrations aim to convey shape characteristics that make designs distinct from one another. Future studies could also test the effects of our illustration style on object recognition tasks (e.g. recognizing ‘chair’ as the object category, rather than differentiating between variations of an object). Further, an ablation study was beyond the scope of our study resources but would be valuable as future work, for example, to study the effect of multi-projection rendering with standard textures.

Our technique relies on a semantically segmented object. Skeleton-based segmentation could be applied, e.g., Livesu et al. [2017], Zhou et al. [2015], Tagliasacchi et al. [2016]. The segmentation could also be obtained with machine-learning approaches, which have recently shown improved performance over traditional approaches [Kalogerakis et al. 2010; Xu et al. 2015].

Finally, other stylizations could be explored, such as exploded diagrams [Li et al. 2008] which might allow for better visibility of occluding parts, abstraction using characteristic curves [Mehra et al. 2009] or geometric primitives to handle more complex geometry, or abstractions such as skeletons to emphasize topological relationships. We also see opportunities for integration with interactive tactile graphics, e.g., combined with interactive audio feedback, dynamic markers [Suzuki et al. 2017], or refreshable tactile graphics

displays where low resolution pin-based stimuli will present new challenges [Brauner 2016].

## 7.2 Conclusion

We have successfully built a general pipeline for the design of tactile illustrations to improve 3D shape understanding for blind individuals. We have also designed a user study to evaluate the effectiveness of the illustrations, which we implemented on a selection of common household objects. We are the first to create a systematic approach that can generate tactile illustrations for objects with complex 3D geometry. We demonstrate promising results, showing significant improvement compared to baseline guidelines for tactile illustration design.

Providing a tool that creates tactile illustrations is integral as a resource for the blind community to allow them to gain access to visual information. For example, our illustrations could be used as a tool to depict products, from furniture in a store catalogue, to art pieces in a museum, to object repositories for 3D printing. An exciting area of future work would be to use our illustration algorithm as a language to build a design tool for the blind. As one of our study participant mentioned: “One of the biggest difficulties blind people have is getting an idea out of my head and into your head.”

## ACKNOWLEDGMENTS

We thank the anonymous reviewers for their feedback; Amber Percy (BU graduate student, MFA Boston Access Consultant, National Braille Press Proofreader) for helpful feedback throughout the project; BU Engineering Product Innovation Center (EPIC) and crew (Kara Mogensen and Caroline) for 3D-printing and laser-cutting resources; VIBUG (Visually Impaired and Blind User Group); Colleen Kim for designing the object modifications for User Study 3; and Forrester Cole for helpful discussions on line drawing and perception. Perkins School for the Blind: Research Librarian Jennifer Arnott and Museum tour-guide Kevin Hartigan provided access to resources, and Executive Director of Perkins Library Kim Charlson and Braille Production Specialist Tanja Milojevic provided help in realizing the study. Museum of Fine Arts Boston: Manager of Accessibility Hannah Goodwin and Accessibility Coordinator Ronit Minchom provided insight, resources and knowledge. We also thank Prof. Sam Ling (BU Visual Neuroscience Lab), who pointed us to the Recognition-by-components theory, Prof. Luis Carvalho (BU Dept. of Mathematics and Statistics) and Prof. James O’Malley (Biomedical Data Science of Dartmouth Institute) for advice on statistical testing, and Prof. Margrit Betke (BU CS Dept.), who highlighted the importance of tactile diagrams in books. Athina also thanks Elizabeth Tremmel (Dartmouth College Multilingual Graduate Student and Postdoctoral Writing Consultant) for writing assistance.

Photograph sources ©Museum of Fine Arts, Boston: Fig. 2(a,b): Pat Lyon at the Forge, John Neagle, 1826-27, Accession Number: 1975.806, Henry H. and Zoe Oliver Sherman Fund; Fig. 2(c,d): Coffee pot, Paul Revere, Jr., 1791, Accession Number: 95.1359, Bequest of Buckminster Brown, M.D. This project is partially supported by the National Science Foundation under Grant No. #1813319, and the Alfred P. Sloan Foundation: Sloan Research Fellowship.

## REFERENCES

2013. Tangent vectors to a 3-D surface normal: A geometric tool to find orthogonal vectors based on the Householder transformation. *Computer-Aided Design* 45, 3 (2013), 683 – 694.
- Maneesh Agrawala, Denis Zorin, and Tamara Munzner. 2000. Artistic Multiprojection Rendering. In *Eurographics Workshop*. Springer-Verlag, London, UK, UK, 125–136.
- Corinna Bauer, Lindsay Yazzolino, Gabriella Hirsch, Zaira Cattaneo, Tomaso Vecchi, and Lotfi Merabet. 2015. Neural correlates associated with superior tactile symmetry perception in the early blind. *Cortex* 63 (02 2015), 104–117.
- Irving Biederman. 1987. Recognition-by-components: a theory of human image understanding. *Psychological review* 94, 2 (1987), 115–147.
- Randolph Blake and Robert Sekuler. 2006. *Perception*. McGraw-Hill, Boston, MA, USA.
- Diane Brauner. 2016. Refreshable Tactile Graphics Display: Making On-Screen Images Accessible! <https://www.perkinslearning.org/technology/posts/refreshable-tactile-graphics-display-making-screen-images-accessible>. Accessed: 2019-08-17.
- Anke M. Brock, Philippe Truillet, Bernard Oriola, Delphine Picard, and Christophe Jouffrais. 2015. Interactivity Improves Usability of Geographic Maps for Visually Impaired People. *Human-Computer Interaction* 30, 2 (2015), 156–194.
- John Brooke. 1996. SUS-A quick and dirty usability scale. 194 (1996), 189–194.
- John Brosz, Faramarz F. Samavati, M. Sheelagh T. Carpendale, and Mario Costa Sousa. 2007. Single Camera Flexible Projection (NPAR ’07). ACM, New York, NY, USA, 33–42.
- Craig Brown and Amy Hurst. 2012. VizTouch: Automatically Generated Tactile Visualizations of Coordinate Spaces (TEI ’12). ACM, New York, NY, USA, 131–138.
- Francis D. K. Ching and Juroszek Steven P. 2019. *Design Drawing*. WILEY, Indianapolis, IN 46256.
- P. Claudet, G. Chougui, and P. Richard. 2000-2008. The Typhlo and Tactus Guide for Children’s Books with tactile illustrations.
- Forrester Cole, K. Sanik, D. DeCarlo, Adam Finkelstein, T. Funkhouser, S. Rusinkiewicz, and M. Singh. 2009. How Well Do Line Drawings Depict Shape? *ACM Trans. Graph.* 28, 3, Article 28 (2009), 9 pages.
- Theresa Cooke, Frank Jäkel, Christian Wallraven, and Heinrich H. Bühlhoff. 2007. Multimodal similarity and categorization of novel, three-dimensional objects. *Neuropsychologia* 45, 3 (2007), 484 – 495. Advances in Multisensory Processes.
- Menena Cottin, Rosana Faria, and Elisa Elisa Amado. 2008. *The Black Book of Colors Hardcover*. Groundwood Books, Toronto, ON M6R 2B7, USA.
- Doug DeCarlo, Adam Finkelstein, Szymon Rusinkiewicz, and Anthony Santella. 2003. Suggestive Contours for Conveying Shape. *ACM Trans. Graph.* 22, 3 (July 2003), 848–855.
- Oliver Deussen, Jörg Hamel, Andreas Raab, Stefan Schlechtweg, and Thomas Strothotte. 1999. An Illustration Technique Using Hardware-based Intersections and Skeletons. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 175–182.
- J. Diepstraten, D. Weiskopf, and Thomas Ertl. 2003. Transparency in Interactive Technical Illustrations. *Computer Graphics Forum* 21 (05 2003), 317 – 325.
- Goker Erdogan, Ilker Yildirim, and Robert A. Jacobs. 2014. Transfer of object shape knowledge across visual and haptic modalities. In *In Proceedings of the 36th Annual Conference of the Cognitive Science Society*.
- J Farley Norman, Hideko F Norman, Anna Marie Clayton, Joann Lianekhammy, and Gina Zielke. 2004. *The visual and haptic perception of natural object shape*. Vol. 66. Springer Berlin Heidelberg, Berlin, Heidelberg, 342–51 pages.
- Stéphane Grabli, Emmanuel Turquin, Frédo Durand, and François Sillion. 2004. Programmable Style for NPR Line Drawing. In *Eurographics Symposium on Rendering*. ACM Press.
- Stéphane Grabli, Emmanuel Turquin, Frédo Durand, and François Sillion. 2008. Freestyle. <http://freestyle.sourceforge.net/>
- Giacomo Handjaras, Emiliano Ricciardi, Andrea Leo, Alessandro Lenci, Luca Cecchetti, Mirco Cosottini, Giovanna Marotta, and Pietro Pietrini. 2016. How concepts are encoded in the human brain: A modality independent, category-based cortical organization of semantic knowledge. 135 (04 2016).
- Lucia Hasty. 1999. Design Principles for tactile graphics. <http://www.tactilegraphics.org/readability.html>
- Morton Heller, Tara M Riddle, Erin K. Fulkerson, Lindsay Wemple, Anne D McClure Walk, Stephanie M Guthrie, Christine Kranz, and Patricia Klaus. 2009. The influence of viewpoint and object detail in blind people when matching pictures to complex objects. *Perception* 38 8 (2009), 1234–50.
- Sergio E. Hernandez and Kenneth E. Barner. 2000. Tactile Imaging Using Watershed-based Image Segmentation (*Assets ’00*). ACM, New York, NY, USA, 26–33.
- Aaron Hertzmann and Denis Zorin. 2000. Illustrating Smooth Surfaces (*SIGGRAPH ’00*). ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 517–526.
- Ian Howard and Robert Allison. 2011. Drawing with divergent perspective, ancient and modern. *Perception* 40 (01 2011), 1017–33.
- Steven Hsiao. 2008. Central mechanisms of tactile shape perception. *Current Opinion in Neurobiology* 18, 4 (2008), 418 – 424. Sensory systems.
- Shoukat Islam, Swapnil Dipankar, Deborah Silver, and Min Chen. 2004. Spatial and Temporal Splitting of Scalar Fields in Volume Graphics. In *Symposium on Volume Visualization and Graphics (VV ’04)*. IEEE Computer Society, Washington, DC, USA,

- 87–94.
- Thomas W. James, G.Keith Humphrey, Joseph S. Gati, Philip Servos, Ravi S. Menon, and Melvyn A. Goodale. 2002. Haptic study of three-dimensional objects activates extrastriate visual areas. *Neuropsychologia* 40, 10 (2002), 1706 – 1714.
- JMS. 2010. True Reverse Perspective. <https://vimeo.com/12518619>
- Tilke Judd, Frédo Durand, and Edward H. Adelson. 2007. Apparent ridges for line drawing. *ACM Trans. Graph.* 26, 3 (2007), 19.
- Evangelos Kalogerakis, Aaron Hertzmann, and Karan Singh. 2010. Learning 3D Mesh Segmentation and Labeling. *ACM Trans. Graph.* 29, 4, Article 102 (July 2010), 12 pages.
- Evangelos Kalogerakis, Derek Nowrouzezahrai, Simon Breslav, and Aaron Hertzmann. 2012. Learning Hatching for Pen-and-ink Illustration of Surfaces. *ACM Trans. Graph.* 31, 1, Article 1 (Feb. 2012), 17 pages.
- J. M. Kennedy. 1993. *Drawing and the Blind: Pictures to Touch*. Yale Univ. Press, New Haven, CT 06520-9040, USA.
- Jeeun Kim, Hyunjooh Oh, and Tom Yeh. 2015. A Study to Empower Children to Design Movable Tactile Pictures for Children with Visual Impairments (*TEI '15*). ACM, New York, NY, USA, 703–708.
- Jeeun Kim, Abigale Stangl, and Tom Yeh. 2014. Using LEGO to Model 3D Tactile Picture Books for Sighted Children for Blind Children (*SUI '14*). ACM, New York, NY, USA, 146–146.
- Jeeun Kim and Tom Yeh. 2015. Toward 3D-Printed Movable Tactile Pictures for Children with Visual Impairments (*CHI '15*). ACM, New York, NY, USA, 2815–2824.
- Roberta L. Klatzky, Jack M. Loomis, Susan J. Lederman, Hiromi Wake, and Naofumi Fujita. 1993. Haptic identification of objects and their depictions. *Perception & Psychophysics* 54 (09 1993), 170–8.
- J. Krikke. 2000. Axonometry: a matter of perspective. *IEEE Computer Graphics and Applications* 20, 4 (July 2000), 7–11.
- Martin Kurze. 1997. Rendering Drawings for Interactive Haptic Perception (*CHI '97*). ACM, New York, NY, USA, 423–430.
- Stephen Lakatos and Lawrence E. Marks. 1999. Haptic form perception: Relative salience of local and global features. *Perception & Psychophysics* 61, 5 (01 Jan 1999), 895–908.
- Manfred Lau, Kapil Dev, Weiqi Shi, Julie Dorsey, and Holly Rushmeier. 2016. Tactile Mesh Saliency. *ACM Trans. Graph.* 35, 4, Article 52 (July 2016), 11 pages.
- S.J. Lederman and Roberta Klatzky. 2009. Haptic Perception: A Tutorial. *Attention, Perception & Psychophysics* 71 (10 2009), 1439–59.
- R Lenth. 2019. emmeans: Estimated marginal means, aka least-squares means. R package version 1.4. 3.01.
- Vincent Levesque. 2005. *Blindness, Technology and Haptics*. Technical Report. Ecole de technologie supérieure.
- Jingyi Li, Son Kim, Joshua A. Miele, Maneesh Agrawala, and Sean Follmer. 2019. Editing Spatial Layouts Through Tactile Templates for People with Visual Impairments (*CHI '19*). ACM, New York, NY, USA, Article 206, 11 pages.
- Wilmot Li, Maneesh Agrawala, Brian Curless, and David Salesin. 2008. Automated Generation of Interactive 3D Exploded View Diagrams. *ACM Trans. Graph.* 27, 3, Article 101 (Aug. 2008), 7 pages.
- Wilmot Li, Lincoln Ritter, Maneesh Agrawala, Brian Curless, and David Salesin. 2007. Interactive Cutaway Illustrations of Complex 3D Models. *ACM Trans. Graph.* 26, 3, Article 31 (July 2007).
- Hsueh-Ti Derek Liu and Alec Jacobson. 2019. Cubic Stylization. *ACM Transactions on Graphics* (2019).
- Marco Livesu, Marco Attene, Giuseppe Patane, and Michela Spagnuolo. 2017. Explicit cylindrical maps for general tubular shapes. *Computer-Aided Design* (05 2017).
- Michael J. McGuffin, Liviu Tancau, and Ravin Balakrishnan. 2003. Using Deformations for Browsing Volumetric Data. In *VIS'03 (VIS '03)*. IEEE Computer Society, Washington, DC, USA, 53.
- N. McLennan, F. Hayden, T. Poppe, and W. Pierce. 1998. The good tactile graphic. Louisville, KY 40206-0085, USA.
- Ravish Mehra, Qingnan Zhou, Jeremy Long, Alla Sheffer, Amy Gooch, and Niloy J. Mitra. 2009. Abstraction of Man-Made Shapes. *ACM Transactions on Graphics* 28, 5 (2009), #137, 1–10.
- Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su. 2019. PartNet: A Large-Scale Benchmark for Fine-Grained and Hierarchical Part-Level 3D Object Understanding. In *CVPR*.
- David Moore. 2011. *Disabled Students in Education: Technology, Transition, and Inclusivity: Technology, Transition, and Inclusivity*. IGI Global, Hershey, PA 17033, USA. 49 – 50 pages.
- A. Moringen, K. Krieger, R. Haschke, and H. Ritter. 2017. Haptic search for complex 3D shapes subject to geometric transformations or partial occlusion. In *2017 IEEE World Haptics Conference (WHC)*. 299–304.
- Braille Authority of North America. 2010. *Guidelines and Standards for Tactile Graphics*. Braille Authority of North America, Pittsburgh, PA, USA.
- Emil Praun, Hugues Hoppe, Matthew Webb, and Adam Finkelstein. 2001. Real-time Hatching (*SIGGRAPH '01*). ACM, New York, NY, USA, 581.
- Paul Ré. 1981. On My Drawings and Paintings: An Extension of the System of Their Classification. *Leonardo* 14, 2 (1981), 106–113.
- Adrian Secord, Jingwan Lu, Adam Finkelstein, Manish Singh, and Andrew Nealen. 2011. Perceptual Models of Viewpoint Preference. *ACM Trans. Graph.* 30, 5, Article 109 (Oct. 2011), 12 pages.
- Abigale Stangl, Jeeun Kim, and Tom Yeh. 2014. 3D Printed Tactile Picture Books for Children with Visual Impairments: A Design Probe (*IDC '14*). ACM, New York, NY, USA, 321–324.
- Abigale J. Stangl. 2019. *Tactile Media Consumption and Production for and by People Who Are Blind and Visually Impaired: A Design Research Investigation*. Ph.D. Dissertation. University of Colorado at Boulder.
- Nisha Sudarsanam, Cindy Grimm, and Karan Singh. 2005. Interactive Manipulation Of Projections With a Curved Perspective. In *Eurographics*.
- Ryo Suzuki, Abigale Stangl, Mark D. Gross, and Tom Yeh. 2017. FluxMarker: Enhancing Tactile Graphics with Dynamic Tactile Markers (*ASSETS '17*). ACM, 190–199.
- Andrea Tagliasacchi, Ibraheem Alhashim, Matt Olson, and Hao Zhang. 2012. Mean Curvature Skeletons. *Comput. Graph. Forum* 31, 5 (Aug. 2012), 1735–1744.
- Andrea Tagliasacchi, Thomas Delame, Michela Spagnuolo, Nina Amenta, and Alexandru Telea. 2016. 3D Skeletons: A State-of-the-Art Report. *Computer Graphics Forum* 35 (05 2016).
- Markus Tatzgern, Denis Kalkofen, and Dieter Schmalstieg. 2010. Compact Explosion Diagrams (*NPAR '10*). ACM, New York, NY, USA, 17–26.
- Anne Theurel, Arnaud Witt, Philippe Claudet, Yvette Hatwell, and Edouard Gentaz. 2013. Tactile Picture Recognition by Early Blind Children: The Effect of Illustration Technique. *Journal of experimental psychology. Applied* 19 (09 2013), 233–40.
- Leanne Thompson and Edward Chronicle. 2006. Beyond visual conventions: Rethinking the design of tactile diagrams. *British Journal of Visual Impairment* 24 (05 2006), 76–82.
- Leanne J. Thompson, Edward P. Chronicle, and Alan F. Collins. 2006. Enhancing 2-D Tactile Picture Design from Knowledge of 3-D Haptic Object Recognition. *European Psychologist - EUR PSYCHOL* 11 (01 2006), 110–118.
- D. V. Vranic, D. Saupe, and J. Richter. 2001. Tools for 3D-object retrieval: Karhunen-Loeve transform and spherical harmonics. In *IEEE Workshop on Multimedia Signal Processing*. 293–298.
- Benjamin Walsh. 2017. Texture, Buttons, Sound and Code: Modal Preference and the Formation of Expert Identities (*FabLearn '17*). ACM, New York, NY, USA, Article 19, 4 pages.
- T. P. Way and K. E. Barner. 1997. Automatic visual to tactile translation. *IEEE Transactions on Rehabilitation Engineering* 5, 1 (March 1997), 95–105.
- Georges Winkenbach and David H. Salesin. 1996. Rendering Parametric Surfaces in Pen and Ink (*SIGGRAPH '96*). ACM, New York, NY, USA, 469–476.
- Kai Xu, Vladimir G. Kim, Qixing Huang, and Evangelos Kalogerakis. 2015. Data-Driven Shape Analysis and Processing. arXiv:cs.GR/1502.06686
- Jeffrey Yau, Sung Soo Kim, Pramodsingh Thakur, and Sliman Bensmaia. 2015. Feeling form: The neural basis of haptic shape perception. *Journal of neurophysiology* 115 (11 2015), jn.00598.2015.
- J. Yu and L. McMillan. 2004. A Framework for Multiperspective Rendering (*EGSR '04*). Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 61–68.
- Johannes Zander, Tobias Isenberg, Stefan Schlechtweg, and Thomas Strothotte. 2004. High Quality Hatching. *Computer Graphics Forum* 23, 3 (2004), 421–430.
- Yang Zhou, Kangxue Yin, Hui Huang, Hao Zhang, Minglun Gong, and Daniel Cohen-Or. 2015. Generalized Cylinder Decomposition. *ACM Trans. Graph.* 34, 6, Article 171 (Oct. 2015), 14 pages.