

# Editing Spatial Layouts through Tactile Templates for People with Visual Impairments

Jingyi Li<sup>1</sup>, Son Kim<sup>2</sup>, Joshua A. Miele<sup>3</sup>, Maneesh Agrawala<sup>1</sup>, Sean Follmer<sup>1</sup>

<sup>1</sup>Stanford University, <sup>2</sup>Vista Center for the Blind and Visually Impaired, <sup>3</sup>The Smith-Kettlewell Eye Research Institute  
{jingyili, maneesh}@cs.stanford.edu, skim@vistacenter.org, jam@ski.org, sfollmer@stanford.edu

## ABSTRACT

Spatial layout is a key component in graphic design. While people who are blind or visually impaired (BVI) can use screen readers or magnifiers to access digital content, these tools fail to fully communicate the content's graphic design information. Through semi-structured interviews and contextual inquiries, we identify the lack of this information and feedback as major challenges in understanding and editing layouts. Guided by these insights and a co-design process with a blind hobbyist web developer, we developed an interactive, multimodal authoring tool that lets blind people understand spatial relationships between elements and modify layout templates. Our tool automatically generates tactile print-outs of a web page's layout, which users overlay on top of a tablet that runs our self-voicing digital design tool. We conclude with design considerations grounded in user feedback for improving the accessibility of spatially encoded information and developing tools for BVI authors.

## CCS CONCEPTS

• **Human-centered computing** → **Accessibility systems and tools**; *Accessibility technologies*;

## KEYWORDS

accessible design tools; accessible web design; layout design; templates; tactile overlays; blindness; visual impairments; multimodal interfaces; accessibility

## ACM Reference Format:


Jingyi Li, Son Kim, Joshua A. Miele, Maneesh Agrawala, Sean Follmer. 2019. Editing Spatial Layouts through Tactile Templates for People with Visual Impairments. In *CHI Conference on Human*

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](mailto:permissions@acm.org).  
*CHI 2019, May 4–9, 2019, Glasgow, Scotland UK*

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300436>



**Figure 1: Our interface leverages interactive tactile sheets to convey web page structure and content. Here, a user double taps on a page element to hear more information.**

*Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA, 11 pages.*  
<https://doi.org/10.1145/3290605.3300436>

## 1 INTRODUCTION

From books to screens, people around the world design and share an abundance of visual media. A rich body of existing work has been dedicated to finding the best way to visually encode, design and present information [28, 30, 34, 47]. Supporting access to this visual language is especially important for people who are blind or visually impaired (BVI) [8]. For example, people who are BVI create slide decks to communicate and present their ideas in the workplace, which is an important employability skill [5]. Many BVI disability advocates create and maintain their own web blogs, showcasing their work not only to BVI peers, but also to sighted colleagues and curious strangers [6, 26].

Currently, people who are BVI use assistive technologies like screen magnifiers, screen readers, and occasionally braille displays to interact with digital content. Screen readers strip out graphical information and present audio descriptions of the page content. Although existing applications—such as PowerPoint or WordPress—allow users to enter text in a pre-defined layout template, BVI authors lack accessible ways of understanding and modifying the template's layout *structure*. Screen readers allow users to tab through PowerPoint's menu options and create widgets in WordPress with their keyboard, but mouse-based interactions such as resizing or repositioning elements are inaccessible. Creating layouts with a screen reader offers no information about nor control over where elements are placed besides their serialized order with regards to other content on the page. While people with low vision who use screen magnifiers can access their mouse, this technology does not provide a sense of the page's entire layout structure at once.

This lack of information about spatial layout makes it difficult for BVI authors to understand how their actions impact the page’s presentation and structure. From a series of contextual inquiries, we observed that their current workflows usually involve entering text into a template and then using a screen reader to speak back that text (whether by presenting their slides or visiting their webpage) to verify their modifications. Many ask sighted collaborators for assistance on further tasks, like rearranging or adding non-textual elements on a page, and “sanity checking” their edits. We see the potential of a tool that supports BVI authors and leverages their tactile and auditory perceptual abilities [48], to foster greater confidence and self-efficacy in not only understanding but also modifying their own layouts.

We carried out an exploration with people who are blind and visually impaired that identified common practices and their limitations in working with layouts. Our findings reveal that BVI people are interested in layout design and want to make modifications to static templates beyond filling in text. Together with co-author Son Kim, a blind Assistive Technology Specialist and hobbyist web designer, we co-designed and prototyped a multimodal tool that makes understanding and editing layouts—specifically web pages—more accessible. Our tool provides BVI authors access to layout structures, augmenting their existing work flows of selecting and filling out layouts. To address the lack of non-visual feedback current technologies provide, our tool automatically generates tactile sheets (which can be laser cut, embossed, or printed on capsule paper [2]) from the bounding boxes of elements on a web page (Figure 1). Similar to prior methods that augment tactile graphics with interactive audio [3, 14, 25], our system uses a tablet to verbalize both the underlying content (e.g., *Project 3 Title*) and meta-content (e.g., *Heading, level 2; text alignment: left*) of the layout. By breaking up what a screen reader would have spoken serially into two modalities, we reduce the amount of information BVI authors have to hold in short term memory.

In this paper, we contribute (1) the results of an initial investigation of how BVI people currently create and edit layouts and the challenges they face, (2) a tool that aids blind people in understanding and editing spatial layouts through multimodal touch interactions, and (3) design considerations for developing accessible authoring tools, such as the trade-offs of different types of spatial representations, developing scaffolding to support users learning domain knowledge, and building systems that amplify the unique perspectives assistive technology users bring to their designs.

## 2 RELATED WORK

Our approach is grounded in prior work in tactile graphics and how to support layout design for sighted people. We draw inspiration from existing cross- and multimodal design tools, and accessible authoring tools.

### Tactile Graphics

There is a wide range of literature [1, 16, 36] specifying best-practice guidelines for creating tactile graphics for people with visual impairments. Since these graphics traditionally have been static—fabricated with a braille embosser or on capsule paper—they are

limited by the amount of information they can display [9]. Commercial products like the Talking Tactile Tablet [25] or the ViewPlus IVEO [14] add interactivity to printed tactile sheets by overlaying them on a touchscreen that plays hand-made audio annotations when touched. To explore tactile graphics, researchers have also used personal fabrication machines. They have employed 3D printing to let blind users access works of art [39], STEM education [11, 20], and traditional WIMP desktop interface elements [4]. Furthermore, researchers have generated static laser cut tactile aids to teach graphic design [31] and provide guidance for using touch screens [21]. Recently, Avila et al. verified that laser cut tactile overlays of document layouts that—like our approach—speak underlying text when touched were as comprehensible as screen readers for text comprehension and navigation [3]. We extend this work with a pipeline to automatically generate such sheets from web pages and an authoring tool for spatial layouts.

Researchers continue to investigate and develop 2D matrices of actuated pins: recent examples include the Holy Braille [33], the HyperBraille [38], and the American Printing House for the Blind’s Graphiti [13]. Although these tactile displays afford dynamic and refreshable graphics, they are prohibitively expensive; for example, the HyperBraille costs \$50,000 USD for a half page size. Focusing on low cost and wide reach, our tool only requires two pieces of hardware: a tablet to run the tool and a machine, like a capsule paper printer or embosser, to fabricate the tactile sheets.

### Layout Design Support

For sighted users, researchers have extensively studied techniques to support layout and graphic design. Design mining of the web has produced explorable design galleries [40] and data driven analyses of website layouts [24]. Both DesignScape [35] and Sketchplore [46] aid novices by offering interactive layout suggestions during the design process. R-ADoMC [19] and Kuhna et al. [23] suggest designs for magazine covers: the former through high and low level descriptors and the latter through image content. Focusing on layout structure, researchers have developed systems to automatically adapt grids to various viewport sizes [18] and make layout structures dynamic and reusable [29]. Our work also targets document structures, but rather than offer design recommendations, we allow BVI authors to understand and modify them.

### Crossmodal and Multimodal Design Tools

Many researchers have looked at enhancing traditionally WIMP-based design tools with multiple modalities—namely, voice and gesture. Speak’n’Sketch [41] aimed to maximize the sketching workspace by abstracting menu options to voice commands, and Crossweaver [43] let designers prototype and test interactive storyboards through speaking and sketching. PixelTone [27], which supported natural language and gestural inputs for image editing, employed preview galleries of available edits to promote learning the system vocabulary. Like PixelTone, our tool facilitates learning through preview galleries; however, we focus not on augmenting visual interfaces but translating them through tactile feedback and audio annotations.

### Authoring Tools for Accessibility

While 3D-printed or laser cut tactile interfaces allow for inexpensive and customizable representations, they rely on pre-fabricated models often made by researchers or experts. To address this issue, researchers have created authoring tools that range from generating tactile overlays of graphics [17] and data visualizations [10], to picture books [22] and annotated 3D models [42]. However, these tools rely on a sighted user to design and produce the tactile artifacts for a blind user, which could limit their independence and ability to more deeply understand the material gained in the creation process. Researchers have thus developed tools that enable blind users to be the creators of their own media. In the domain of audio editing, The Moose [37] enabled blind users to splice audio by acting as a powered mouse with force feedback, while the Haptic Wave [44] translated digital waveforms into the haptic channel. Facade [15] employs crowdsourcing to let blind users generate overlays for household appliances like microwaves, and Taylor et al. present a system to create interactive 3D printed tactile maps [45]. Like these systems, our tool allows blind users to create for themselves; we focus specifically on the domain of comprehending and iterating upon layout designs.

### 3 FORMATIVE STUDIES

To better understand how BVI people engage in layout design, we conducted semi-structured interviews with two low vision and five blind participants (Table 1). Additionally, we draw on the experiences of the second and third authors, who are both blind. We returned to three participants to conduct contextual inquiries where we observed them creating and editing layouts for presentation slides and websites, the domains our interviewees expressed were the most challenging.

#### Procedure

We recruited participants by contacting local community organizations, email lists, and snowball sampling. We conducted an hour long semi-structured interview with each participant that highlighted the challenges in accessing spatial information without tactile aids. We revisited the offices of P2, 3, and 6 for contextual inquiries where we observed them creating a website (P2) and/or slide deck (P2, P3, P6). Data consisted of audio and video recordings, field notes, and presentations and websites participants sent us. We used a modified grounded theory [12] approach to code our interview findings. Our code categories fit into two main themes: participants' current practices of layout design and the challenges they encountered. We used this framing to help organize insights from the contextual inquiries, and present our final code categories as the subsections below.

#### Findings: Current Practices

*Workflows.* Our participants created layouts mainly in the context of their work—everyone had experience making slide decks, and P2 and P7 maintained their own websites. Participants began with content in mind and would either outline it in a word processor to then paste into a layout template (P2), or write their content directly into the template (P3, 6), and then speak or magnify the content to verify their actions. All participants made their slide decks in

PowerPoint, except for P1 who used Google Slides. P6, who is low vision, used memorized combinations of mouse movements to access menu options in PowerPoint, while blind participants used their keyboards to tab through and select the options. After using their keyboard to select a text box identified by its placeholder text, blind participants would type their own content and not make further changes themselves. For websites, P2, who knew how to program, edited HTML and CSS files directly. P7 used WordPress for its interface that enabled selecting and creating widgets (like images or post comments) from the keyboard. This is consistent with the second author's own experiences of editing websites, who also used to directly edit HTML files, but prefers the higher level interactions of WordPress.

*Choosing Templates.* Participants chose templates from higher level descriptors, if they chose a template besides the default one at all. For example, P7 searched for WordPress themes that used the "accessible" keyword, which mirrors the second author's experiences. Because screen readers do not convey graphical information like the typography or color palette of templates, some participants relied on sighted colleagues to choose a "pretty enough" or "socially acceptable" theme. P2 mentioned, "I don't feel confident in experimenting in choosing new layouts, because I would need someone to explain to me what it looked like every time. I just go with the default—that people have already described—because I know it." P3 and P6 sent us PowerPoint presentations they had created without external help, and both used the default template of black Calibri text on a white background.

*Modifications.* During our contextual inquiries, blind participants did not modify templates beyond entering text. They mentioned when they wanted to add new non-textual content (such as images), they would ask someone sighted. Outside of content-based modifications, participants expressed a desire to control graphical components of layouts that they believed would impact sighted peoples' judgments of their sites. This was aligned with the experiences of the second author, who would ask sighted colleagues to help choose new color palettes for his websites, because he wants them to look appealing and professional, and not just "another website which uses the same old template." Low vision participants modified themes to be even more accessible. P7, who runs a low vision advocacy blog, dramatically increased the font size in her theme, and P6 placed colorful widgets on her slides as place markers so she would know where she was during her presentation, since the text was too small for her to read.

#### Findings: Challenges

*Lack of feedback.* The biggest challenge blind participants faced in understanding and editing layouts was that screen readers did not provide information about the spatial relationships encoded in layout elements, such as position and size. Instead, they experienced layouts through serialized audio descriptions of textual content. P5 mentioned that his modifications to templates was "absolutely a verification problem." P3 echoed these sentiments, saying, "Even if I have the ability to make a change, I don't know how it actually affected the layout. I just trust the computer did it, or ask someone."

ID	Gender	Age	Vision level and onset	Desktop accessibility features	Layouts made & frequency
1	Male	22	Born blind with light perception	Screen reader	Word documents (weekly), Slides (on occasion)
2	Male	23	Born blind	Screen reader	Official documents (monthly), Websites, Slides (on occasion)
3	Female	61	Born blind with light perception, total vision loss at 50	Screen reader	Slides (on occasion)
4	Female	51	Born blind	Screen reader	Slides (thrice a week), Word documents (weekly)
5	Male	52	Blind at age 1	Screen reader	Slides (on occasion)
6	Female	28	20/450 at age 17	Screen magnification & text-to-speech	Slides, Graphs (weekly), Websites, Fliers (on occasion)
7	Female	21	Low vision at age 3	Screen magnification & reader, high contrast display, blue light filter	Websites (weekly), Slides (on occasion)

**Table 1: Participant demographics. Participants 1-5 are blind, and 6 & 7 are low vision.**

Screen readers sometimes inadequately communicated the results of blind participants’ editing actions. For instance, we noticed text overflowing out of its container in a website P2 had designed and edited, because he accidentally deleted a closing tag in the HTML and his screen reader did not report the error. Similarly, when the second author added a video inside WordPress’s content editor, it was automatically resized to fit the full page width. He immediately asked, “How does it look, visually?” and assumed the video was small and centered on the page. Without proper feedback, the mental representations BVI authors held of the layout were mismatched with the actual layout.

*Lack of exposure to design standards.* Because they have mainly experienced web pages as a spoken ordered list of page content, congenitally blind participants did not have opportunities to build up their knowledge of standard layouts and design trends. P1 shared, “Before I try to make my own PowerPoints pretty, I need to ask what people think is pretty first.” In the instance of the fit to width video, the second author was concerned visitors would be off put by the large video player, asking if it “looked abnormal,” even though fit to width media have become standard in current design trends. Furthermore, the kinds of information screen readers exposed, such as font names, were unhelpful if participants didn’t know what the font looked like in the first place. P2 wished for a “design score” that evaluated the page’s adherence to graphic design principles, and P6 specifically asked sighted colleagues if her slides were visually appealing or not. These mirrored the experiences of the third author, who mentioned, “What I care about is, does the font choice make the page look childish, or cramped?”

*High load on short-term memory.* Screen readers translate the parallel experience of vision into a serial experience of speech. While this serialization is adequate for accessing individual elements, it makes exposing and understanding relationships between elements challenging due to all the information BVI people must hold in their short-term memory. The third author described his experiences with interpreting dense modern web pages through a screen reader as “playing chess in [his] head.” Participants with low vision who used screen magnifiers also faced the issue of being unable to understand the gestalt of their layouts, because big picture relationships are lost when zooming in.

*Sighted collaborators.* Several participants had experience making layouts in collaborative contexts (i.e., P6 was an editor of an academic magazine and P2 was a start up co-founder and had to make slide decks to raise venture capital). While they trusted their sighted collaborators to graphically design the content they provided, their collaborators sometimes lacked the proper domain knowledge to sufficiently visually communicate content only they understood, such as making graphs (P6) or technical diagrams (P2). P7 also mentioned the social strife that could result from creative differences with collaborators. “When you ask someone else to do it, you’re reliant on their judgment and understanding and tastes, which usually aren’t yours. And sometimes they take it personally, and then you’re in a social mess.”

*Confidence.* None of our blind participants could move, align, or resize content in their layout with a screen reader without external help. We observed low vision participants using menu options and buttons to align elements; P6 said she avoided direct manipulation with the mouse for she feared her interactions inaccurate. This lack of confidence that they would not “mess up” the layout prevented BVI users from editing content in the first place. P2 said, “I lack the ability to make last minute changes—I would never make a change to my presentation without having a sighted person double check it, not even if it’s as small as changing a word. I just don’t want to screw it up.”

### Design Imperatives

In summary, like sighted people, BVI people utilize templates in creating layouts. However, their current assistive technologies restrict them to filling in content and making high level modifications through PowerPoint or WordPress menus, like changing color or type. Because they lack exposure to and feedback on layout elements, BVI users express a desire to understand what “good” layout designs are. They often juggle many elements in their short term memory and have mental models differing from what they are designing. BVI authors rely on sighted collaborators for assistance, which they mentioned reduced their confidence in designing independently.

Based on these findings, we developed several design guidelines for our tool to help BVI authors in designing and understanding layouts. We want to (1) leverage their existing workflows of starting

from templates (i.e., by adapting our tool to work with existing templates); (2) provide feedback on and the ability to make spatial edits (i.e., through interactive tactile sheets); (3) present content-layout relationships in multiple modalities to avoid high cognitive loads (i.e., by offloading some of the screen reader’s audio to the tactile channel); and (4) support learning of unfamiliar layout designs and concepts (i.e., by exposing options in the context of editing).

#### 4 AN ACCESSIBLE LAYOUT DESIGN TOOL

Our tool enables users understand layout structures and make modifications in both page content and layout. It uses automatically generated tactile sheets overlaid on a tablet that speaks and edits HTML templates. We support two different layout representations of these templates—*pixel* and *inflated*—that respectively conform to the original web page design and tactile graphics guidelines.

While our formative studies revealed insights both in domains of slide and web design, we chose to build a web-editing tool for a few reasons. First, the web is a platform with a rich tradition of broadening access to information, and web layouts have complex issues that are widely applicable to other layout tasks (which we expand on in the discussion.) Second, web pages have an underlying textual declarative representation (HTML/CSS) for visual layouts, while this representation is hidden in other domains like closed-source slide deck software. Our interface allows users to edit layouts by manipulating this underlying representation, which generalizes to a platform more widely used than creating our own specialized software representations. Finally, we wanted to leverage building a tool considering our co-designer’s interests and domain expertise.

Although our formative study participants had a variety of visual impairments, to address the issues brought forth by screen readers, we frame our tool specifically for people who are blind. We imagine our tool to also be useful for people with low vision, but acknowledge the wide spectrum of visual impairments and push against a “one-size-fits-all” solution.

##### Co-Design Process

We closely designed our tool with co-author Son Kim, an Assistive Technology specialist who is also a blind hobbyist web designer. We iterated upon three design probes before arriving at our current implementation. For the first probe, we generated three sets of tactile sheets by hand and encoded different page elements—like paragraphs versus images—with different textures. We chose to use tactile sheets because blind people use two hands with multiple fingers while exploring tactile layouts [32], and the sheets afford multiple finger and whole hand exploration. Our goals were to verify that tactile sheets would lead to an adequate mental map of spatial relationships as others have demonstrated [3], and to understand how BVI designers might choose layouts with this additional understanding over solely high-level tags. Son recreated layouts on a felt board to demonstrate his understanding, and discussed with the researchers where perceptual challenges occurred, which were usually around closely-spaced elements.

In the second probe, we generated tactile sheets of two website templates and mounted them on a tablet that provided audio annotations. We presented the tool without editing features to focus on feedback on how it helped with layout comprehension. Son

took about 10 minutes to understand each layout, and suggested creating a tactile sheet representation that exposed the web page’s underlying grid structure and to help him more quickly identify where elements were. We incorporated this functionality in revealing container relationships in the tool’s “inflated” mode.

Finally, our goal in the last probe was to identify a set of editing actions to support. We presented five web page templates to modify. Son demonstrated what kinds of interactions and edits he wanted to make, which we carried out with through Wizard of Oz manipulations of the underlying website by changing the DOM in real time and refabricating new tactile sheets to reflect his edits.

Three themes emerged through the co-design sessions. First, legibility difficulties in the pixel-perfect layout representation reemphasized the importance of tactile graphics guidelines such as leaving a minimum 1/8th inch between elements [16]. Second, with spatial relationship knowledge, Son was able to assign semantic meaning to the elements—for instance, he chose a layout in the first probe for a storefront because the large space allocated for an image made it “showy and screamed, ‘Buy this now!’” Third, breakdowns occurred due to a lack of past exposure and contextual knowledge of components (such as unfamiliarity with z-ordering or pagination), mirroring our formative study findings.

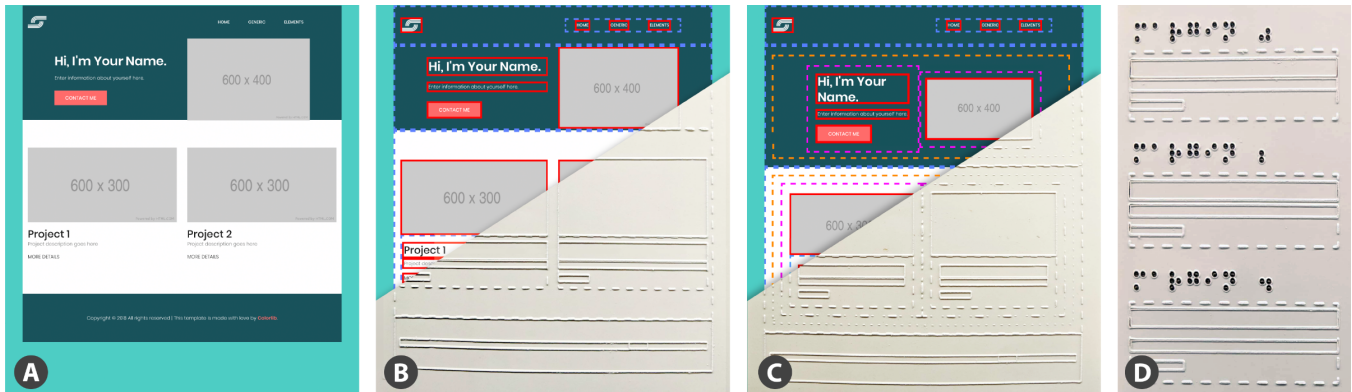
##### Understanding and Editing Layouts with Our Tool

We describe our tool with a hypothetical design scenario. Ben, a blind student, wishes to make a personal portfolio website to showcase his projects before applying to jobs. He starts by searching for templates using tags, such as “accessible” and “portfolio,” and downloads a prospective HTML file, akin to his current workflow, following from design imperative (1).


*Providing access to spatial components and relationships.* Ben loads the first HTML file (Figure 2A) into our tool on his tablet and selects the *Print* button in the toolbar (Figure 3). This downloads an SVG file that is the *pixel* representation of the layout—that is, bounding box outlines of page elements (Figure 2B). He prints this SVG file on capsule paper and runs the paper through a capsule paper printer to generate raised lines. He affixes this new tactile sheet to his tablet.

Ben now explores this web page by running his fingers over the capsule paper. Ben double taps on a large element on the left side of the layout, and he hears information about the element: *Element: image. Alternate text: Project 1 image placeholder.* He then taps below that, *Element: Heading level 2. Content: Project 1. Text align: left.* This interaction follows from design imperative (2), and splitting layout information across modalities—structure to touch and meta information to audio—supports design imperative (3). After exploring the other two templates in this fashion, Ben chooses the first layout because he likes that the layout isn’t too cluttered and that it has partitioned sections for both his personal information and portfolio.

*Editing content and layout.* First, Ben fills in his content text by double tapping the *Edit* button in the tool bar to toggle on edit mode. In edit mode, Ben first selects the element he wants to modify—here, the Heading level 1 to replace “Your Name” with “Ben.” He double taps the *Back* and *Next* buttons to scroll through a menu of operations, such as editing, deleting, swapping, or duplicating



**Figure 2:** A sample HTML template (A) and our tool’s automatically generated pixel (B) and inflated (C) representations, fabricated by laser cutting capsule paper. We also generate small multiples (D); here, we show several levels of bottom margin between the project header and description. Note the outlines of the bounding boxes are for figure purposes only.



**Figure 3:** A menu appears at the top of the design tool. Users double tap the menu options to change the tool’s modality and cycle through editing options.

the element, and traverses through the menu hierarchy: *Selected Heading level 1* → *Edit element* → *Edit text*. After Ben enters text using a keyboard connected to the tablet, he also fills in the images he had planned for his website by specifying their URLs.

With the content filled in, Ben now wants to modify the layout structure—he has three projects he wants to showcase, but the template came with space for two. He toggles the print mode to generate an *inflated* representation of the layout that exposes elements’ containers (Figure 2C). This representation adds padding to container elements to comply with tactile graphic guidelines [16], including a minimum 1/8th distance between lines, and exposes the underlying grid structure of the page. He selects an HTML column (Figure 4.1), chooses the *duplicate* action, and then selects the new container to duplicate the column in—in this case, its row (Figure 4.2). Our tool supports adaptive layouts and asks Ben if he would like to keep the current column sizes, or shrink them to keep all the elements on the same row; he selects the latter (Figure 4, after). Now there is a mismatch in representation between the underlying layout and tactile sheet, so Ben reprints and replaces the overlay, which he keeps track of by a braille label at the top.

Ben now notices that his project heading is too closely spaced to its paragraph description. He wants to add a bottom margin to the element, but isn’t sure about what value would best match the existing layout design. With the *Edit margin* action selected, Ben double taps the *Print* button to generate a tactile sheet of a preview gallery. This gallery contains multiple copies of his project text description container (with the heading, paragraph text, and “More Details” link) with different levels of bottom margin applied to the heading (Figure 2D). Following design imperative (4), the preview gallery aims to expose potentially unfamiliar concepts situated in the context of the page. Ben applies a wider margin to

the heading, and then selects the heading and uses the *Copy style* action to propagate the changes to the other two project headings. Finally, Ben prints a pixel representation of his layout to verify how his edits would appear on most screens.

### Implementation

Our design tool relies on manipulating the web page’s Document Object Model (DOM). To generate tactile sheets, our tool traverses the DOM, gets an object’s bounding box coordinates, creates an SVG rectangle. Generated tactile sheets have four different line patterns: solid for elements (Figure 4D), dashed for general containers (4C), dot-dashed for columns (4B), and dotted for rows (4A). Elements are considered containers if their inner text matches the sum of their children’s inner text. Notably, we cannot simply just consider non-leaf nodes containers to avoid common false positives, such as marking a paragraph a container if it has a link child. To generate an inflated layout, we add 1/8th inch of padding to any container element whose padding is currently less than that. Because we use responsive templates, increasing padding does not horizontally overflow elements at the tradeoff of increased element heights. To generate in-context preview galleries of small multiples, we extract and modify the selected element in its parent container, applying all available levels of padding, margin, alignment, or resizes.

Once users download the SVG, they can then fabricate a tactile sheet. Users calibrate sheets to their devices by specifying their screen width and a height offset. We used a laser cutter to fabricate sheets, which takes an average of 10 to 120 seconds per sheet depending on the complexity. However, sheets may be generated on other 2-axis CNC machines, like an embosser, or with a capsule paper printer, which requires a short fixed time (about 15 seconds) regardless of layout complexity. Once the sheets are fabricated, users affix them to a tablet; we attach ours to an iPad Pro by securing them under a beveled case.

Actions our tool currently supports include creating, deleting, duplicating, and swapping content elements or their containers, which we refer to collectively as layout components. Users can edit layout components by editing their text content or tags (e.g., changing a paragraph to a heading), resizing them, aligning them




Figure 4: In the inflated representation, we expose HTML rows (A), columns (B), other containers (C), and content elements (D). Before: To duplicate a column, first users select the element to be duplicated (the column) (1) and then in the container to duplicate it in (the row) (2). After: Users have the option to adaptively resize the new columns in the resultant row.

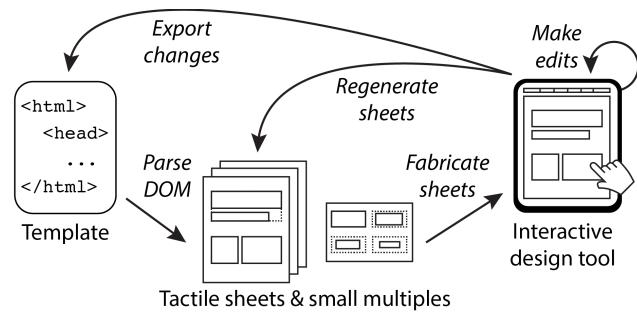


Figure 5: Our system takes an HTML template as input and allows users to iteratively edit templates and reprint and replace them as tactile sheets. Once users are done editing, they can export the edited HTML file.

within their container, text aligning them, editing their padding or margin, or copying their style to another layout component. We chose these actions based on the results of our co-design process.

Our tool is a Node.js web application. It takes an HTML file as input, and outputs both pixel and inflated SVG representations of our tactile sheets and the edited HTML file (Figure 5). We employ the ResponsiveVoice.js library for speech feedback, Hammer.js to handle multitouch interactions, and Bootstrap as our templating framework.

## 5 RESULTS

We tested our tool against popular website and web-based slide templates and show some of the generated tactile sheets in Figure 6. In addition to developing our tool with an experienced web designer, we verified its usability and received positive informal feedback from non-domain experts.

### Informal Evaluation

Since our tool was co-designed with an experienced web designer, we wanted to understand how the tool helps lower the threshold for layout comprehension and design with novices. We recruited two blind participants with no experience in creating layouts for informal feedback. Neither participated in the formative studies nor considered themselves familiar with tactile graphics. P1 lost his vision nine years ago and has been using a screen reader on his desktop computer since; P2 lost her vision at eight months and did not regularly use a desktop computer, although she began using a smart phone three months prior.

Participants underwent one hour of training with our tool and one hour of a design task. During training, the proctor guided them through all of the editing features provided. For the design task, participants first explored and chose from four Bootstrap templates for personal portfolio websites, and then filled in content and modified the layout.

Both participants were able to successfully identify and recall layout elements, as well as populate them with content and make structural layout modifications. A template took from 15 to 20 minutes to fully comprehend. Participants initially had trouble decoding the different dashed line representations for containers, but learned after hearing the container type spoken a few times. Both participants swapped tactile sheets when their edits resulted in a large change in the page structure—like inserting an image that was longer than its placeholder space—but were fine with relying on audio feedback for actions like alignment and swapping elements that were the same size. P2 preferred the inflated layout representations and said the wider spaces between elements were more helpful to understanding “which boxes were the walls of the room, and which ones were the furniture.” Both participants found the design galleries helpful in understanding unfamiliar concepts like padding. “Because I added another picture, I would want to decrease the padding between them to loosen up the page,” said P1.

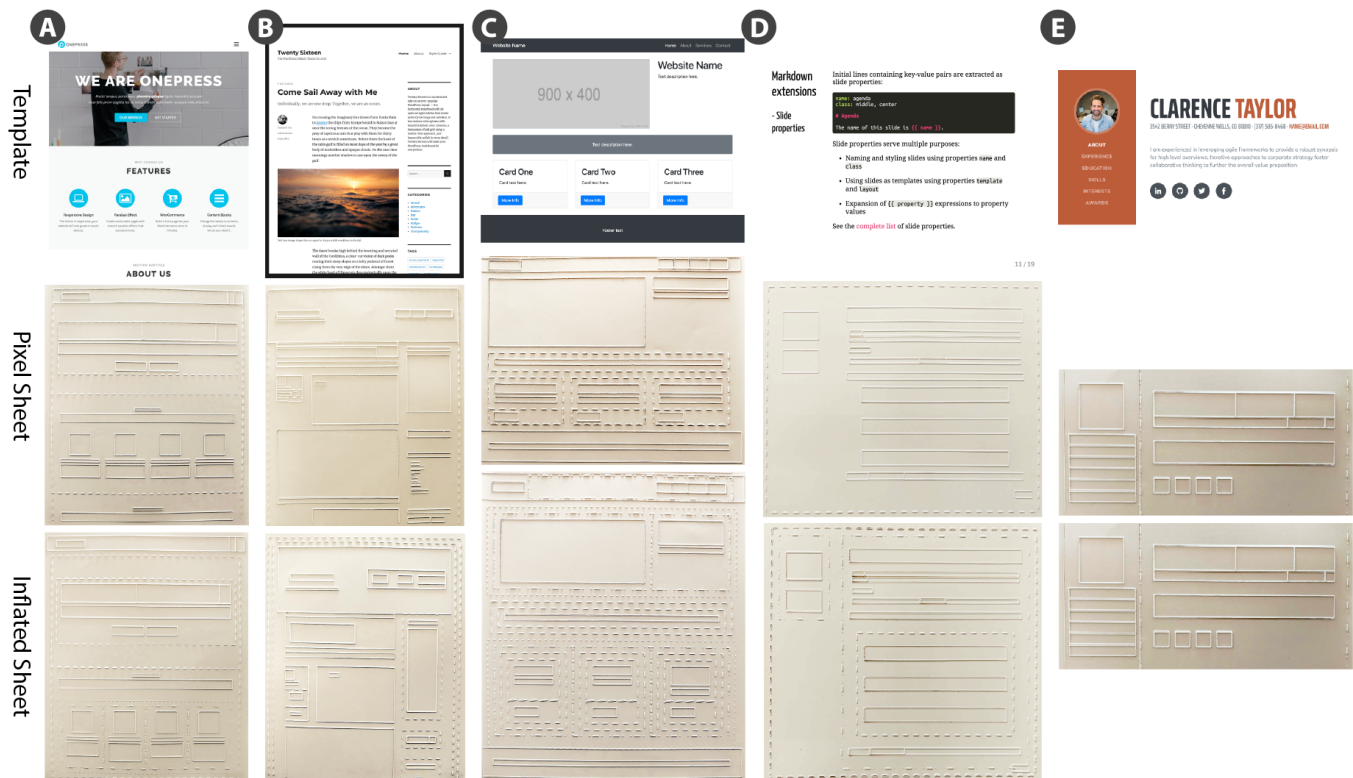
Despite their lack of layout experience, both participants found the tool useful. “It’s a lot more detailed than I expected, but I like that level of specific control. With more time to familiarize myself prior work goes from being an impossibility to being possible,” said P1. “I thought editing a website would be difficult because I didn’t know anything about web pages and sites, but I actually enjoy editing. I feel more control when I can place in images. I think I understand a lot more about web pages now and am more confident in my ability to make one,” said P2.

## 6 DISCUSSION AND FUTURE WORK

### Limitations

People who are blind or visually impaired have diverse abilities and experiences. It is important to note that all our formative study interviewees had (or was working towards) a college degree and were experienced users of accessible desktop technology.

Our informal evaluation also had a small sample size. Future work should focus on a longitudinal “in the wild” deployment with blind web designers.



**Figure 6: Sample templates taken from popular themes and their corresponding pixel and inflated tactile representations. A and B are Wordpress themes, C and E are Bootstrap themes, and D is a sample slide from Remark.js, an HTML-based slide library. Note that since template E did not have containers that needed to be inflated, its representations are equivalent.**

While our tool exposes layout structures, it does not communicate other design aspects like typography or color. Translating these elements requires more research on choosing the appropriate descriptors that facilitate their understanding for congenitally blind people. Our tool affords *iterative* feedback and editing, but the time it takes to fabricate and replace the tactile sheets prevents an *interactive* workflow. A refreshable tactile display [13, 38] could bypass this issue.

Our algorithm to generate tactile sheets sometimes falsely exposes span elements that are used for stylistic markers (e.g., changing a font color) but irrelevant to spatial layout design (for example, the last name in Figure 6E.) Since it relies on the bounding box coordinates of elements, elements with bounding boxes that stretch to fill to their parent but with sparse content appear to take up more space than they actually do (for example, the footer text in Figure 6C.) These cases could be handled with a vision-based approach.

### Design Considerations of Tools for BVI Designers

Through our co-design sessions and informal evaluations, several themes emerged which we believe to be valuable to other researchers developing authoring tools for BVI users: the trade-offs between pixel-perfect layout representations and tactile graphics, how to support a wide range of abilities and backgrounds in developing domain-specific tools, and building tools that recognize

the valuable first-hand perspectives users of assistive technologies bring to their designs.

*Pixel-Perfect Views and Tactile Graphic Guidelines.* When generating tactile sheets, there is a tension between representing the structure of a web page as a direct pixel translation of how it appears in the browser and modifying it to conform to tactile graphics standards. Many child nodes share the same borders as their parent containers, or are adjacent to other elements. These overlapping borders are difficult to differentiate tacitly—all users confused negative spaces (like the space between a navigation bar link and the end of the bar) as small layout elements. Adding more space between element borders addresses this issue; however, increasing spacing leads to either overflow (with increased margins) or taller, “squished” content (with increased padding). Our tool chooses the latter to preserve column structures central to modern web templates as to avoid unnecessary layout re-comprehension. One possible intervention is to only fabricate content elements and present the container information in audio—for example, elements could also speak the names of their parent containers. Another option is to allow zooming in to the layout and add spacing at a finer granularity for less disruption. Although more research needs to be done to evaluate the trade-offs of different tactile representations for layout design for BVI users, our findings can inform the design of tools built on refreshable tactile displays, which would be



particularly well suited for zooming and quickly toggling between layout representations.

*Scaffolding for Novice Designers.* As evidenced by our informal evaluation, people have a wide range of visual impairments and familiarity with domain-specific concepts. Our co-designer noticed patterns of organizational structures throughout our sessions, such as template elements being arranged side-by-side in the same row, and not top-to-bottom as he assumed was the case since he used the top and bottom arrow keys to access elements with his screen reader. However, those without web design experience struggled with understanding web-specific jargon, like “div.” In addition to generating preview galleries that explain some terms in context, our tool could come with a binder of simple examples that emphasize different principles, or include training methods grounded in best practices of teachers of the visually impaired.

We designed our tool to facilitate the existing practices of users, allowing them to select and modify individual elements at the same level as their screen readers. However, web design novices found understanding our lower-level representations of a layout’s underlying grid structure challenging. Due to concerns about the accuracy of direct manipulation from the formative studies and past work [7], we avoided pixel-based manipulations and computationally abstract direct editing of CSS or HTML. High-level design tools have supported sighted novice designers [27, 35], and we believe the same applies for users who are BVI. Our tool could offer global heuristic-based design “sanity checks,” or execute edits following high-level commands like “make this layout less cluttered.”

*Supporting Edits for Accessibility.* We observed users making edits with accessibility in mind—for instance, Son swapped the image and text columns so screen reader users would be able to have an image description first. While we have developed a tool for blind people that represents how sighted individuals perceive websites, we do not yet support how blind people could design websites for other screen reader users. For example, a third layout representation could expose a screen reader’s reading order of page elements, and allow modifications to this order for a better user experience.

## 7 CONCLUSION

We present the challenges people who are blind or visually impaired face in understanding and editing spatial layouts, and address these challenges by developing a multimodal tool to help blind people understand and edit web pages by using tactile templates. Drawing from our experiences co-designing the tool with a blind hobbyist web developer and feedback from two blind web design novices, we point out the tensions in representing spatial layouts with tactile sheets and accommodating for a wide range of expertise. We believe BVI users bring valuable accessibility-oriented perspectives to web design and push for greater inclusion of their voices and insights.

## ACKNOWLEDGMENTS

We wish to thank the formative study and informal evaluation participants for their valuable time and insights, as well as Alexa Siu, Elyse Chase, Eric Gonzalez, Evan Strasnick, Jennifer Jacobs, and Niamh Murphy for their helpful conversations and assistance. This work is supported by an NSF GRFP Fellowship #2017245378.

## REFERENCES

- [1] Nacny Amick and Jane Corcoran. 1997. Guidelines for Design of Tactile Graphics. <http://www.aph.org/research/guides/>. Accessed: 2018-09-06.
- [2] Sona K Andrews. 1985. The use of capsule paper in producing tactual maps. *Journal of Visual Impairment & Blindness* (1985).
- [3] Mauro Avila, Francisco Kiss, Ismael Rodriguez, Albrecht Schmidt, and Tonja Machulla. 2018. Tactile Sheets: Using Engraved Paper Overlays to Facilitate Access to a Digital Document’s Layout and Logical Structure. In *Proceedings of the 11th PErvasive Technologies Related to Assistive Environments Conference (PETRA ’18)*. ACM, New York, NY, USA, 165–169. <https://doi.org/10.1145/3197768.3201530>
- [4] Mark S. Baldwin, Gillian R. Hayes, Oliver L. Haimson, Jennifer Mankoff, and Scott E. Hudson. 2017. The Tangible Desktop: A Multimodal Approach to Nonvisual Computing. *ACM Trans. Access. Comput.* 10, 3, Article 9 (Aug. 2017), 28 pages. <https://doi.org/10.1145/3075222>
- [5] Edward C. Bell and Natalia M. Mino. 2015. Employment outcomes for blind and visually impaired adults. (2015).
- [6] Edward C. Bell and Natalia M. Mino. 2016. Blogging when you’re blind or visually impaired. <https://yourdolphin.com/news?id=223>. Accessed: 2018-08-30.
- [7] Jens Bornschein, Denise Bornschein, and Gerhard Weber. 2018. Comparing Computer-Based Drawing Methods for Blind People with Real-Time Tactile Feedback. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI ’18)*. ACM, New York, NY, USA, Article 115, 13 pages. <https://doi.org/10.1145/3173574.3173689>
- [8] Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P. Bigham. 2013. Visual Challenges in the Everyday Lives of Blind People. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’13)*. ACM, New York, NY, USA, 2117–2126. <https://doi.org/10.1145/2470654.2481291>
- [9] 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.
- [10] Craig Brown and Amy Hurst. 2012. VizTouch: Automatically Generated Tactile Visualizations of Coordinate Spaces. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction (TEI ’12)*. ACM, New York, NY, USA, 131–138. <https://doi.org/10.1145/2148131.2148160>
- [11] Erin Buehler, Shaun K. Kane, and Amy Hurst. 2014. ABC and 3D: Opportunities and Obstacles to 3D Printing in Special Education Environments. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility (ASSETS ’14)*. ACM, New York, NY, USA, 107–114. <https://doi.org/10.1145/2661334.2661365>
- [12] Kathy Charmaz and Linda Liska Belgrave. 2007. Grounded theory. *The Blackwell encyclopedia of sociology* (2007).
- [13] American Printing House for the Blind. 2016. American Printing House for the Blind and Orbit Research Announce the World’s First Affordable Refreshable Tactile Graphics Display. <https://www.aph.org/pr/aph-and-orbit-research-announce-the-worlds-first-affordable-refreshable-tactile-graphics-display/>
- [14] John A Gardner, Vladimir Bulatov, and Holly Stowell. 2005. The View-Plus IVEO technology for universally usable graphical information. In *Proceedings of the 2005 CSUN International Conference on Technology and People with Disabilities*.
- [15] Anhong Guo, Jeeun Kim, Xiang ‘Anthony’ Chen, Tom Yeh, Scott E. Hudson, Jennifer Mankoff, and Jeffrey P. Bigham. 2017. Facade: Auto-generating Tactile Interfaces to Appliances. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI ’17)*. ACM, New York, NY, USA, 5826–5838. <https://doi.org/10.1145/>

- 3025453.3025845
- [16] Lucia Hasty. 2018. Guidelines for Design of Tactile Graphics. <http://www.tactilegraphics.org/index.html>. Accessed: 2018-09-06.
- [17] Liang He, Zijian Wan, Leah Findlater, and Jon E. Froehlich. 2017. TacTILE: A Preliminary Toolchain for Creating Accessible Graphics with 3D-Printed Overlays and Auditory Annotations. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '17)*. ACM, New York, NY, USA, 397–398. <https://doi.org/10.1145/3132525.3134818>
- [18] Charles Jacobs, Wilmot Li, Evan Schrier, David Bargeron, and David Salesin. 2003. Adaptive Grid-based Document Layout. *ACM Trans. Graph.* 22, 3 (July 2003), 838–847. <https://doi.org/10.1145/882262.882353>
- [19] Ali Jahanian, Jerry Liu, Qian Lin, Daniel Tretter, Eamonn O'Brien-Strain, Seungyon Claire Lee, Nic Lyons, and Jan Allebach. 2013. Recommendation System for Automatic Design of Magazine Covers. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces (IUI '13)*. ACM, New York, NY, USA, 95–106. <https://doi.org/10.1145/2449396.2449411>
- [20] Shaun K. Kane and Jeffrey P. Bigham. 2014. Tracking @Stemxcomet: Teaching Programming to Blind Students via 3D Printing, Crisis Management, and Twitter. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE '14)*. ACM, New York, NY, USA, 247–252. <https://doi.org/10.1145/2538862.2538975>
- [21] Shaun K. Kane, Meredith Ringel Morris, and Jacob O. Wobbrock. 2013. Touchplates: Low-cost Tactile Overlays for Visually Impaired Touch Screen Users. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '13)*. ACM, New York, NY, USA, Article 22, 8 pages. <https://doi.org/10.1145/2513383.2513442>
- [22] Jeeun Kim and Tom Yeh. 2015. Toward 3D-Printed Movable Tactile Pictures for Children with Visual Impairments. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2815–2824. <https://doi.org/10.1145/2702123.2702144>
- [23] Mikko Kuhna, Ida-Maria Kivelä, and Pirkko Oittinen. 2012. Semi-automated Magazine Layout Using Content-based Image Features. In *Proceedings of the 20th ACM International Conference on Multimedia (MM '12)*. ACM, New York, NY, USA, 379–388. <https://doi.org/10.1145/2393347.2393403>
- [24] Ranjitha Kumar, Arvind Satyanarayan, Cesar Torres, Maxine Lim, Salman Ahmad, Scott R. Klemmer, and Jerry O. Talton. 2013. Webzeitgeist: Design Mining the Web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 3083–3092. <https://doi.org/10.1145/2470654.2466420>
- [25] Steven Landau and Lesley Wells. 2003. Merging tactile sensory input and audio data by means of the Talking Tactile Tablet. In *Proceedings of EuroHaptics*, Vol. 3. 414–418.
- [26] Veronica Lewis. 2018. Veronica With Four Eyes. <https://veroniiiica.com/>. Accessed: 2018-08-30.
- [27] Jason Linder, Gierad Laput, Mira Dontcheva, Gregg Wilensky, Walter Chang, Aseem Agarwala, and Eytan Adar. 2013. PixelTone: A Multimodal Interface for Image Editing. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*. ACM, New York, NY, USA, 2829–2830. <https://doi.org/10.1145/2468356.2479533>
- [28] Ellen Lupton. 2004. Thinking with type. *Critical Guide for Designers, Writers, Editors & Students* (2004).
- [29] Nolwenn Maudet, Ghita Jalal, Philip Tchernavskij, Michel Beaudouin-Lafon, and Wendy E. Mackay. 2017. Beyond Grids: Interactive Graphical Substrates to Structure Digital Layout. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 5053–5064. <https://doi.org/10.1145/3025453.3025718>
- [30] David McCandless. 2012. *Information is beautiful*. Collins London.
- [31] Samantha McDonald, Joshua Dutterer, Ali Abdolrahmani, Shaun K. Kane, and Amy Hurst. 2014. Tactile Aids for Visually Impaired Graphical Design Education. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility (ASSETS '14)*. ACM, New York, NY, USA, 275–276. <https://doi.org/10.1145/2661334.2661392>
- [32] Valerie S Morash, Allison E Connell Pensky, Steven TW Tseng, and Joshua A Miele. 2014. Effects of using multiple hands and fingers on haptic performance in individuals who are blind. *Perception* 43, 6 (2014), 569–588.
- [33] Valerie S Morash, Alexander Russomanno, R Brent Gillespie, and Sile O'Modhrain. 2017. Evaluating Approaches to Rendering Braille Text on a High-Density Pin Display. *IEEE transactions on haptics* (2017).
- [34] Josef Müller-Brockmann. 1996. Grid systems. *Alemanha, Editora Braun* (1996).
- [35] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2015. DesignScape: Design with Interactive Layout Suggestions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1221–1224. <https://doi.org/10.1145/2702123.2702149>
- [36] Braille Authority of North America. 2010. Guidelines and Standards for Tactile Graphics. <http://www.brailleauthority.org/tg/>. Accessed: 2018-09-06.
- [37] M Sile O'ÁŽModhrain and Brent Gillespie. 1997. The moose: A haptic user interface for blind persons. In *Proc. Third WWW6 Conference*.
- [38] Denise Prescher, Jens Bornschein, Wiebke Köhlmann, and Gerhard Weber. 2018. Touching graphical applications: bimanual tactile interaction on the HyperBraille pin-matrix display. *Universal Access in the Information Society* 17, 2 (2018), 391–409.
- [39] Andreas Reichinger, Helena Garcia Carrizosa, Joanna Wood, Svenja Schröder, Christian Löw, Laura Rosalia Luidolt, Maria Schimkowitzsch, Anton Fuhrmann, Stefan Maierhofer, and Werner Purgathofer. 2018. Pictures in Your Mind: Using Interactive Gesture-Controlled Reliefs to Explore Art. *ACM Trans. Access. Comput.* 11, 1, Article 2 (March 2018), 39 pages. <https://doi.org/10.1145/3155286>
- [40] Daniel Ritchie, Ankita Arvind Kejriwal, and Scott R. Klemmer. 2011. D.Tour: Style-based Exploration of Design Example Galleries. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 165–174. <https://doi.org/10.1145/2047196.2047216>
- [41] Jana Sedivy and Hilary Johnson. 1999. Supporting Creative Work Tasks: The Potential of Multimodal Tools to Support Sketching. In *Proceedings of the 3rd Conference on Creativity & #38; Cognition (C&#38;C '99)*. ACM, New York, NY, USA, 42–49. <https://doi.org/10.1145/317561.317571>
- [42] Lei Shi, Yuhang Zhao, and Shiru Azenkot. 2017. Markit and Talkit: A Low-Barrier Toolkit to Augment 3D Printed Models with Audio Annotations. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17)*. ACM, New York, NY, USA, 493–506. <https://doi.org/10.1145/3126594.3126650>
- [43] Anoop K. Sinha and James A. Landay. 2003. Capturing User Tests in a Multimodal, Multidevice Informal Prototyping Tool. In *Proceedings of the 5th International Conference on Multimodal Interfaces (ICMI '03)*. ACM, New York, NY, USA, 117–124. <https://doi.org/10.1145/958432.958457>
- [44] Atau Tanaka and Adam Parkinson. 2016. Haptic Wave: A Cross-Modal Interface for Visually Impaired Audio Producers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2150–2161. <https://doi.org/10.1145/2858036.2858304>

- [45] Brandon Taylor, Anind Dey, Dan Siewiorek, and Asim Smailagic. 2016. Customizable 3D Printed Tactile Maps As Interactive Overlays. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '16)*. ACM, New York, NY, USA, 71–79. <https://doi.org/10.1145/2982142.2982167>
- [46] Kashyap Todi, Daryl Weir, and Antti Oulasvirta. 2016. Sketchplore: Sketch and Explore with a Layout Optimiser. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 543–555. <https://doi.org/10.1145/2901790.2901817>
- [47] Jan Tschichold. 1998. *The new typography: A handbook for modern designers*. Vol. 8. Univ of California Press.
- [48] Jacob O. Wobbrock, Shaun K. Kane, Krzysztof Z. Gajos, Susumu Harada, and Jon Froehlich. 2011. Ability-Based Design: Concept, Principles and Examples. *ACM Trans. Access. Comput.* 3, 3, Article 9 (April 2011), 27 pages. <https://doi.org/10.1145/1952383.1952384>