TactualPlot: Spatializing Data as Sound using Sensory Substitution for Touchscreen Accessibility

Pramod Chundury, Yasmin Reyazuddin, J. Bern Jordan, Jonathan Lazar, and Niklas Elmqvist

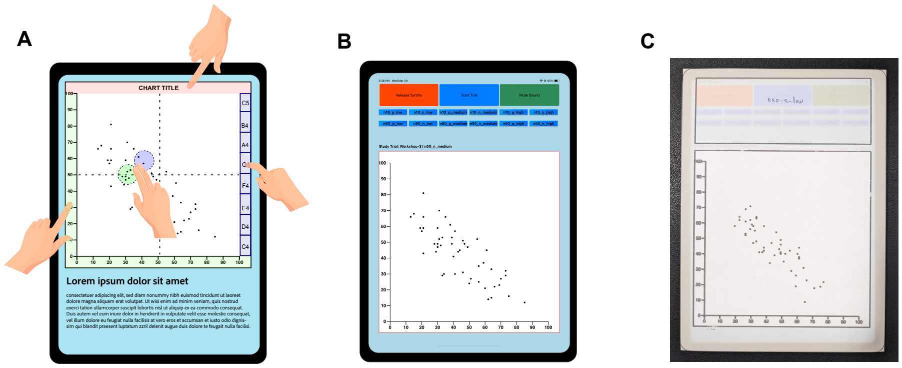

Fig. 1: **Designing the TactualPlot technique.** (A) Illustration of the TactualPlot technique with its continuous and discrete touch actions and regions. (B) Our implementation of the technique that was evaluated in formative design session 3 and design review sessions. (C) Tactile scatterplot implemented using American Thermoform swell touch paper, and aligned to dimensions of the D3.js visualization in the TactualPlot system. Tactile scatterplots were overlaid on an Apple iPad to collect touch information.

Abstract—Tactile graphics are one of the best ways for a blind person to perceive a chart using touch, but their fabrication is often costly, time-consuming, and does not lend itself to dynamic exploration. Refreshable haptic displays tend to be expensive and thus unavailable to most blind individuals. We propose TACTUALPLOT, an approach to *sensory substitution* where touch interaction yields auditory (sonified) feedback. The technique relies on embodied cognition for spatial awareness—i.e., individuals can perceive 2D touch locations of their fingers with reference to other 2D locations such as the relative locations of other fingers or chart characteristics that are visualized on touchscreens. Combining touch and sound in this way yields a scalable data exploration method for scatterplots where the data density under the user's fingertips is sampled. The sample regions can optionally be scaled based on how quickly the user moves their hand. Our development of TactualPlot was informed by formative design sessions with a blind collaborator, whose practice while using tactile scatterplots caused us to expand the technique for multiple fingers. We present results from an evaluation comparing our TactualPlot interaction technique to tactile graphics printed on swell touch paper.

Index Terms—Accessibility, sonification, multimodal interaction, crossmodal interaction, visualization.

1 Introduction

The sense of touch is critical for most blind individuals. Bereft of full use of their vision, blind people often use their hands to explore unfamiliar objects or their white cane to learn about the world surrounding them. Building on this idea, *data physicalization* [29] creates tangible representations of data to enable a blind person to feel rather than see the physical shape of a bar chart [51], 3D landscape [36], or node-link graphs [16]. However, data physicalizations that are cheap and accessible, such as thermoform paper and 3D printing, are static and require significant time to produce. More advanced data physicalization techniques that enable interactive feedback—such as shape-changing displays [3], haptic touch displays [45], and refreshable Braille displays—rely on specialized technology that is expensive and therefore not readily available to an often underemployed blind audience with limited purchasing power. Furthermore, the accessibility field is rife with "silver bullet" technologies that have been abandoned by over-optimistic inventors [35], making investing in experimental devices both costly and risky. Additionally, collaboration between sighted and blind individuals in data-rich social contexts such as work and school might still be challenging due to technology that places sighted and blind individuals in two different silos. While these sophisticated technologies may make data accessible for blind individuals, adoption at scale is not guaranteed because of the aforementioned limitations.

What is readily available, however, are touchscreen devices. Modern smartphones have had a near-revolutionary impact on blind individuals, putting screen readers into virtually everyone's pocket [13]. While still costly, smartphones have significant utility for both personal and professional use and thus enjoys widespread adoption—more than 46% [1]—among blind people, and the number is likely higher for gainfully employed knowledge workers. What if we could design a method for sensing complex datasets on smartphones and tablets using the same touch-based exploration that blind people routinely use with physical objects but without requiring specialized haptic technology?

In this paper, we propose to achieve this by producing audio rather than haptic output in response to touch exploration. This approach is based on *sensory substitution* [4, 14], a common technique from assistive technologies where stimulus for one sensory system is replaced with stimulus for another, such as a screen reader verbalizing written text on a screen with synthesized speech. Our use of sensory substitu-

- *Pramod Chundury, J. Bern Jordan, and Jonathan Lazar are with University of Maryland, College Park in College Park, MD, USA. E-mail: {pchundur, jbjordan, jlazar}@umd.edu.*
- *Yasmin Reyazuddin is with the National Federation of the Blind in Baltimore, MD, USA. E-mail: yasmin81065@gmail.com.*
- *Niklas Elmqvist is with Aarhus University in Aarhus, Denmark; the work was conducted at University of Maryland, College Park, MD, United States. E-mail: elm@cs.au.dk.*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

tion is crossmodal, a special case of *multimodal substitution* [15, 37] utilizing multiple sensory channels, because the approach explicitly "crosses" touch input with sound output. The experience of crossmodal substitution is similar to *synesthesia*—a perceptual phenomenon in which stimulation of one sense leads to involuntary experiences in a second sensory system. In our approach we cross the sense of touch with sound, an approach similar to producing sound when seeing colors [41, 59]. To explore and validate this idea, we propose TACTUALPLOT, a crossmodal substitution technique for scatterplots that enables a person to "touch data" spatialized in two dimensions on a smartphone touch screen, similar to a visual 2D scatterplot. With TactualPlot, a blind user can use their fingertips to explore the shape of the data, receiving a dynamically changing audio tone in response that conveys the density under the user's touch. Unlike many other accessible data representations, TactualPlot was designed for scalability, with thousands of data points in mind.

We designed TactualPlot in a user-centered manner by engaging our collaborator and co-author, who is blind, in a series of formative design sessions. Our observations yielded several design revisions, including the need for multiple touch interaction and axis manipulation. Our participatory design sessions helped us understand that it is important to develop a touch-based, non-visual graphical perception technique that could be generalized beyond scatterplots. We then conducted an in-depth expert review with two blind professionals to assess the utility of TactualPlot, and—by extension—the validity of the crossmodal substitution concept. During the hour-long sessions, our experts used TactualPlot to explore data and answer questions about their findings.

The main contributions of our work include: (*i*) the design and implementation of the TactualPlot prototype for exploring scatterplots using sensory substitution where tactile feedback is sonified; (*ii*) results from an in-depth and longitudinal formative design process involving our blind collaborator and coauthor; and (*iii*) findings from an expert review involving two blind professionals using TactualPlot for multidimensional data.

2 BACKGROUND

Visualization has long mostly ignored the fact that visual representations are not accessible for blind or low vision (BLV) individuals [18], but this is now changing. However, designing accessible visualization is fraught with complexity, both in terms of technical challenges (e.g., how to represent complex and large-scale data primarily using sound and touch) as well as social (adoption, maintenance, and training) and economical (high cost to an often underemployed user population) barriers [35]. Here we review the literature on accessible visualization, focusing mostly on the aforementioned technical challenges.

2.1 Sensory and Multimodal Substitution

Sensory substitution refers to the use of one sense to provide information normally provided by another sense [4], and is a common approach for assistive technology in accessibility [14]. Audiobooks are a prime example of assistive technology using sensory substitution; originally invented to help blind individuals enjoy reading books, their widespread adoption also with sighted individuals showcases a common phenomenon in accessibility: the "curb-cut effect," where improvements for one population of users end up benefiting many. *Multimodal substitution*, on the other hand, involves combining different sensory modalities to create a richer experience [15, 37]. For instance, combining touch, sound, and smell can provide a more complete representation of an environment for a blind user.

Sensory and multimodal substitution are commonly used also for accessible visualization [13]. This is primarily done through sound and tactile feedback; we describe these research efforts in detail below. However, researchers have also explored the potential of using smell to make visualization accessible for blind individuals [5, 43].

2.2 Chart Accessibility

While the accessibility community has long worked towards accessible graphics, it is only recently that the visualization community has begun to pay attention [17, 18, 20, 31, 50]. Given the understandable reluctance among the blind community to adopt untested and poorly maintained technology [35], one strategy is to target trusted assistive technology.

Screen readers may be one such opportunity. Alternate texts (alt-texts) are machine readable descriptions associated with images on the web, and are commonly verbalized by screen readers to help blind individuals access image content via screen readers. Jung et al. [30] propose a comprehensive set of guidelines for writing alternative text descriptions for visualizations to cater to the diverse needs of blind individuals. However, the adoption of alt-text on the internet is poor even for regular images [22], let alone chart images. To deal with these situations, Al-Zaidy and Giles presented an algorithm that uses computer vision and OCR techniques to automatically extract data from bar charts with an accuracy of over 90% [2]. Similarly, Choi et al. [11] proposed an approach to reverse-engineer several types of rasterized charts to make the data accessible for blind users.

Another approach is designing visualizations to be highly-compatible with a screen reader. Zong et al. [61] worked with blind collaborators to design visualizations whose structure, navigation, and descriptive content are optimized for rich screen reader experiences. The VoxLens [48] integrates with the screen reader to convey a multimodal approach to visualization accessibility, providing voice commands, summarizing data, and sonifying details on demand. Most recently, Thompson et al. [52] worked over a period of five months to develop a chart accessibility engine combining a screen reader, data sonification, and descriptive content generation for web-based charts.

2.3 Sonifying Data

In a state-of-art report written for the U.S. National Science Foundation in 1997, Kramer et al. [34] define *data sonification* as "*the use of non-speech audio to convey information*" (p. 4), further qualifying it as "*the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation*" (p. 4). Sonification is a form of an *auditory display* [33]; other forms (not necessarily mutually exclusive with sonification) include audification (directly converting large-scale data to the audible domain, auditory icons (short and self-contained sounds representing discrete events), and verbalization (synthesized speech conveying data).

The ICAD community has been the epicenter of auditory display research for the last 30 years, but not all auditory displays are concerned with conveying data; some approaches are mostly artistic in natures, whereas others focusing on conveying realistic soundscapes rather than abstract data. Nevertheless, many sonification efforts can be applicable for such abstract datasets; the *Sonification Handbook* [25] surveys the state of the art in the field.

One of the early sonification approaches was the iSonic system [60], where spatialized geographic data is conveyed using sound to support a blind user navigating and querying a map using physical key mappings. Our work in this paper is heavily inspired by iSonic, but draws on the widespread adoption of smartphones where the touchscreen becomes the equivalent to the physical keys. The web-based chart library Highcharts has recently begun distributing an accessibility tool called the Sonification Studio [9], which enables robust and flexible data exploration using sound. Wang et al. [58] performed a study to rank audio channels in the sonification of data, confirming that pitch is optimal for encode data, but that tappings and length can be effective for specific tasks or data types. Potluri et al. [44] recently proposed PSST, a toolkit enabling blind developers to author accessible data displays. Finally, Hoque et al. [27] present a study on how the use of natural sounds can enable blending multiple data channels in parallel for increasing the sensory bandwidth of the sonification. Holloway et al. [26] use both sonification and speech to make infographics accessible. While we use sonification in our approach, we do not sonify the data directly, but rather we sonify the spatial attribute of data in a scatterplot.

2.4 Touching Data

Standard touch-based assistive technology [14] include the ubiquitous white cane, Braille text that can be read using the fingertips, and tactile maps [26], whereas Braille keyboards enable blind individuals to generate text. However, for representing data, the options are limited.

Static tactile graphics are mostly made through 3D printing, thermal printing (where lines and shapes are raised when heat is applied), or embossing (which press designs into paper).

Digital tactile representations have the benefit of being able to refresh dynamically, but are often specialized or costly, or both. Nevertheless, such refreshable technologies have been proven effective in conveying data from visual representations into touch-perceivable equivalents, as evidenced by studies on maps [26] and bar charts [51]. Guinness et al. [24] found that using miniature robots to convey data through tactile feedback was more effective for target acquisition than using sound. Such data physicalization [29], shape-changing displays [3], and haptic touch displays [45] could well present workable solutions. However, with the exception of the long-awaited Dynamic Tactile Device (DTD)[1] being developed by HumanWare and American Printing House for the Blind, most of these advanced devices are costly research prototypes, and thus are not widely available to the general blind population. Another option may be the commercially available ultrahaptics display [10], which generates mid-air tactile sensations using ultrasound, but the device still provides a fairly low resolution.

The nearly ubiquitous smartphone [1] may be a better solution since it incorporates both audio output and a touch surface. A recent paper explores the use of touch for exploring 2D visualizations to yield sonified 3D sound output [39]. However, their evaluation uses six blindfolded sighted participants, which is questionable because sighted participants lack the lived experiences of blind individuals, and thus this is not an ecologically valid approach [18]. While our work is based on a similar idea, our approach to multi-touch interaction using a sampling region is more robust and was iteratively developed in a participatory design process with a blind collaborator and tested with blind experts.

3 Approach: Crossmodal Substitution

Sensory substitution is the practice of replacing one sensory system for another when producing output for some perceptual task [4, 14]. It is a common approach for assistive technologies because it enables replacing a sensory system inaccessible to a person with disabilities with a sensory system that remains accessible. For example, a screen reader uses sensory substitution to replace text on a screen—which requires vision to perceive—with synthesized speech—which uses hearing. This allows a blind person to read digital documents, webpages, and articles.

The notion of *simultaneous crossmodal sensory substitution* (SCSS), or just *crossmodal substitution*, replaces in real time the output for one sensory system that is produced by the user's interaction with output for another sensory system. A specialized form of *multimodal sensory substitution* [15,37], which uses multiple sensory channels for feedback, crossmodal substitution only makes sense in a digital space where the normally rich physical interactions of the real world do not hold. For example, manipulating a 3D object using a standard mouse typically yields none of the touch, feel, and heft of manipulating an artifact in the real world. However, visual feedback still yields a facsimile of the sensory output inherent with handling such a physical object; it may move, rotate, and even have inertia similar to the real object.

Crossmodal substitution is essentially a specialization of *direct manipulation* [49], and is clearly a fundamental (and unremarkable) aspect of graphical user interfaces. However, it has an important additional affordance when applied to assistive technologies for people with disabilities: it can be used to separate the means of manipulating the world (or the digital system) from the sensory output normally produced by this interaction. More specifically, it could be used to, for example, let a deaf person tap a sound source in a video—such as a car, person, or airplane—and receive a textual description of the nature of the sound, or to let a blind person tap a webpage to get an audio representation of the content (which is the idea explored in this paper). Another use may be to couple an eye tracker with a screen reader, producing a system that automatically verbalizes whatever text the user's eyes alight on.

In this paper, we specifically focus on the idea of letting users interact with a touch screen containing spatialized data to produce a sonification and verbalization of the data. However, as indicated above, we think there is much additional potential in the crossmodal substitution concept that is worth exploring in the future. For example, modern smartphones equipped with vibration motors can be leveraged to include vibrotactile feedback in addition to sound output. In the next section, we describe our approach to crossmodal substitution that utilizes touch, as our input modality; and audio as the output modality. Broadly, we propose mapping discrete touches (tapping) to speech, and continuous touches (swiping and panning) to non-speech sounds. Our technique was iteratively developed through participatory design sessions with a blind collaborator; and that helped us scope the design space to explore only sound and touch at this point in time.

[1] https://www.aph.org/dtd-fact-sheet/

4 The TactualPlot Technique

TactualPlot (Figure 1) is a crossmodal substitution technique for exploring multidimensional data using touch input to produce sound output. The technique spatializes the data by projecting data items to the physical space defined by the extents of the touch screen, similar to a 2D scatterplot for visual representations. Users can then explore these spatialized points by touching them on the screen, similar to how you can explore a rough 3D surface using your fingers. However, instead of providing haptic feedback through your fingertips, the feedback is crossmodally redirected to sound: the data is sonified using a continuous audio tone. The audio playback is derived from sampling the data points under the user's touch and modulating the pitch according to the data density. In this way, TactualPlot can interactively sonify hundreds of data points even on a small touch surface and give a blind user an understanding of the data distribution through progressive exploration.

In this section, we summarize and list the basic principles behind TactualPlot's spatial mapping and interaction models. These principles were finalized based on early discussions with our research team, and informed by our iterative design process described in Section 5. We also describe various interaction methods for touching data as well as the edge of the screen, and the resulting verbalizations and sonifications for representing data. Please note that we did not implement and assess all the features described in this section during our formative design sessions. We recommend carefully implementing these features as the interplay of sound, touch, and data abstraction can be cognitively taxing for our users. The features we implemented during the design sessions, and demonstrated during the expert review are marked with a circle and check mark icon: ✔. Features that we designed, but did not implement are denoted by an empty solid circle: ●. We chose to implement the features that were crucial for the visualization tasks in Table 1.

4.1 Spatial Mapping

Effective data representations for blind users are often based on spatializing—rather than visualizing—data [18]. TactualPlot maps data to space using a direct projection from a data dimension to a geometric dimension. The technique maps data to the horizontal (x) and vertical (y) axes; in other words, it is a 2D spatial representation, similar to a visual scatterplot. Most commonly, the projection from data space to visual space is linear, although logarithmic mappings can be useful.

The TactualPlot spatial mapping is designed to consume the full screen of the device, i.e. its entire length and width. This gives a blind user a direct mapping from the physical display itself—as well as from proprioceptive feedback from their own body—to the symbolic data dimensions in the dataset being represented.

Note that while TactualPlot shares many similarities with a 2D scatterplot, the marks sonified using the technique can really only convey two scalar data dimensions (the two Cartesian axes) and one categorical data dimension (the data item class). Scatterplots, on the other hand, can use the color, shape, and size of marks to convey additional data dimensions.

4.2 Interaction Model

TactualPlot's interactions are designed according to a fundamental rule: *discrete* actions, such as tapping or double-tapping the screen, yield *discrete* feedback (verbalized information), whereas *continuous* actions, such as dragging and pinching on the display, yield *continuous* feedback (continuous audio sonifying the underlying data). In practice, this means that tapping the left margin of the screen will cause TactualPlot to read the name of the data dimension mapped to the vertical axis as well as its current minimum and maximum values, whereas dragging your finger along the axis will read out its current axis values. We discuss these interactions below.

The TactualPlot direct spatial mapping to a physical device facilitates perceiving physical space as data space. This allows, for example, a blind user to slide their finger from left to right and understand that the price of the houses they encounter on the display is increasing.

4.3 Continuously Touching Data

Reading data off a TactualPlot display is achieved by touching the display, which will aggregate the data points contained within the *sampling area* centered on the user's finger and translate the points into audio feedback. Using a circular sampling area [19] rather than a point makes it easier for the user to find all of the data on the plot.

The sampling area can be either set to a specific size (roughly corresponding to the physical size of a typical fingertip) or a size controlled by the user. Another approach is to use a dynamic size for the sampling area that changes based on the movement speed of the user's finger: the faster the finger moves, the larger the sampling area. This can improve the likelihood of sampling all the points on the chart canvas. This enables fast and ballistic movements to form a rough overview, e.g. quickly scribbling with your finger across the display, followed by a slower and more deliberate tracing of the data to get the details.

Every time the user's touch moves on the display, points that fall within the sampling area are re-aggregated. If the points are of different classes, such as a dataset of real estate containing both condominiums, single-family houses, and farms, they can be partitioned into separate groups. The number of points in each group is then normalized globally for the entire chart and the calculated for the current density.

4.4 Data Sonification

Data in a sampling area being sonified using TactualPlot is conveyed using pitch to represent data density; the higher the density, the higher the pitch. Pitch rather than volume has been shown in past work on sonification to be the optimal channel to convey quantitative values [13].

In keeping with the continuous nature of the drag interaction, the sound generated in TactualPlot is also continuous. As recommended [47], we wanted a pleasant sound, and so so implemented a flute like sound. The tone persists as long as the user's finger is touching the display, and it modulates smoothly in pitch as the user moves their finger and the data density under the sampling area changes.

If the underlying dataset has multiple classes, such as different types of real estate, different voices (such as instruments) can be played simultaneously. Again, to maintain consistency with the continuous interaction, individual tones must not be segmented or split. In other words, this requires polyphonic sound generation where multiple tones are played at the same time. Another approach is to use natural sounds that have been shown to blend well together [27].

Finally, TactualPlot can use spatial sound output—stereo ✓ or full 3D binaural audio—to redundantly encode the position of the touch point on the physical device. Alternatively, different stereo channels can be used to divide multiple data classes; for example, playing the current density of condos in one ear, and the density of single-family houses in another. However, blind users commonly use only one headphone while interacting with a mobile device [13], especially in situations where they need to listen to conversations in their surroundings, so techniques that require both ears are not always practical.

- ✓ *Drag - Sonify Sampled Data:* The basic interaction of TactualPlot; the data in the sampling region under the user's fingertips is sampled and the density conveyed using pitch.
- ✓ *Tap - Drill Down:* Retrieve a list of the data values in the sampling region under the user's fingertip (see Section 4.6).
- ● *Two Finger Tap/Pinch - Zoom:* Change the zoom level or center-point of the viewport (see Section 4.6).

4.5 Edge Interactions

The edges of the TactualPlot display are significant and interactions are different from when interacting with the main area of the display. Typically you would define the 10% of the outer parts of the display as the edge. We define four different edge regions with specific affordances:

- ● **Top edge - Title:** Interactions for the entire chart.
  - ● *Tap - Chart Title:* The title of the chart is verbalized.
  - ● *Swipe - close chart:* The TactualPlot display is closed down, yielding a verbal status message.
- ✓ **Left edge - Vertical axis:** Interacting with the y axis.
  - ● *Double Tap - Axis title:* Vertical axis title is verbalized.
  - ✓ *Drag - Axis ticks:* Tick marks are indicated using an earcon.
  - ✓ *Tap - Axis ticks:* Tick mark value is verbalized.
- ✓ **Bottom edge - Horizontal axis:** Interacting with the x axis.
  - ● *Double Tap - Axis title:* Horizontal axis title is verbalized.
  - ✓ *Drag - Axis ticks:* Tick marks are indicated using an earcon.
  - ✓ *Tap - Axis ticks:* Tick mark value is verbalized.
- ● **Right edge - Legend:** Interacting with audio legends.
  - ● *Tap - Data density legend:* The audio class tone of the sonification is played and their corresponding point values are verbalized. In our prototype, tapping directly on the points conveys pitch mappings to the point density.

Overall, all state changes in the TactualPlot interface are accompanied with a text-to-speech status message giving the user verbal feedback for their interaction.

4.6 Zooming and Details-on-Demand

In addition to continuously touching the data, the main touch surface (and not the edge) of the display supports several additional interactions:

- ● **Pinching (two fingers):** Users can zoom into a region of the data display by pinching, thereby changing the data extents on the vertical and horizontal axes. Discrete zoom values in multiples of magnification are verbalized as the display is zoomed in and out. The display cannot be zoomed out past $1x$ magnification, where all of the data is contained within the viewport. Leaving the display untouched for some time (30 to 60 seconds) will revert the magnification back to $1x$, which is also announced verbally. This is done to avoid a blind user returning to a device after some time and not remembering that the display is zoomed in and having no easy way to understand this from the interface.
- ✓ **Tapping (one finger):** Users can drill down into the data to get details-on-demand by tapping on a region. The data density inside the sampling area around that touch point will be verbalized. If there is none, a brief sound or no sound can play. If there is more than a predefined number of points (5 or more), the summary statistics of the points can be verbalized.
- ● **Zoom level (discrete):** Double tapping on the screen can be used to convey spatial information (geometric zoom), or the designer can choose to enable zooming in the data.

4.7 Beyond Scatterplots

The TactualPlot technique was designed specifically to be an interactive sonification of 2D scatterplots, and uses many of the same metaphors and interactions as a visual scatterplot would. However, these principles could no doubt also be applied to other types of visualizations to yield comparable interactive sonifications. Such work is outside the scope of this paper and is left for future research.

5 FORMATIVE DESIGN ASSESSMENT: TACTUALPLOT

TactualPlot was improved using a user-centered, participatory approach. We derived the original design (Section 4) through in-depth discussions in our research team, where one of our collaborators (and coauthor of this paper) is a blind individual with long experience in assistive technologies and human-computer interaction.

Tactile graphics are often used by teachers of students with visual impairments (TVSI) to teach graphical perception in educational contexts. Blind individuals have varying levels of expertise in perceiving representations such as Braille and tactile graphics. Nevertheless, tactile representations are ideal for conveying spatial awareness, especially since multiple fingers can provide parallel channels for tactile information. Therefore, we decided to base our TactualPlot technique on tactile chart exploration with a crossmodal representation using sound.

5.1 Design Probes

To help inform our design process, we built two design probes [21, 57] as low-fidelity prototypes: (1) tactile graphics implemented using swell touch paper, and (2) a prototype web-based app running on smartphones and tablets. Our goal was to let design lessons and findings for the tactile graphics scatterplots inform the design of the mobile application.

Tactile graphics. The tactile graphs were implemented in consultation with our university's assistive technology lab. We performed two iterations of the scatterplot design; the first session enabled us to fixate on printing parameters such as dot sizes, aspect ratio, and data density, whereas the second yielded graphics suitable for both this formative design as well as the subsequent design sessions and user study (Section 6). We describe each session and its outcomes below.

We used swell touch paper made by American Thermoform in standard U.S. Letter size (8.5 × 11 inch; 215.9 mm × 279.4 mm). Scatterplots were printed using a standard laser printer and then "fused" using an American Thermoform Swell Form Machine, which causes ink to swell. It is possible to achieve different heights for the graphical elements by varying the (1) saturation of the ink, and (2) the temperature setting during the fusing process. In our study, we use a single height across the entire tactile graphic. We chose the point diameter in the tactile graphics and the digital TactualPlot system to the base dot size of the North American Braille—1.44 mm; 0.057 inches. This may help retain familiarity among Braille literate blind users, and the size is proven to be easily perceivable to touch.

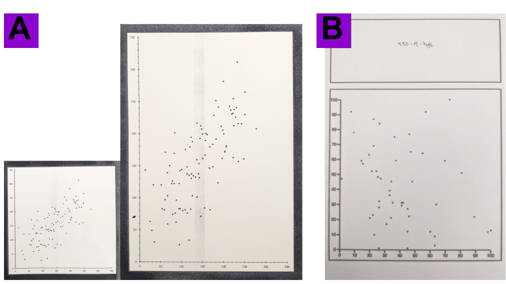

Fig. 2: **Tactile graphics.** (A) Square-shaped scatterplot with uniform axis length, and a rectangle-shaped plot with longer vertical axis used in Design Session 1. (B) Final square-shaped uniform scatterplot used in Design Session 3 for assessing data visualization tasks.

TactualPlot implementation. We implemented TactualPlot as a web-based application for mobile platforms (smartphones and tablets) using basic HTML, CSS, and JavaScript technologies as well as Apache Cordova.[2] Our implementation generates sound using ToneJS,[3] a JavaScript toolkit based on the Web Audio W3C API.[4] Finally, touch interaction is detected and handled using the Touch Web API that can handle multi-touch interactions.[5]

[2]https://cordova.apache.org/
[3]https://tonejs.github.io/
[4]https://www.w3.org/TR/webaudio/
[5]https://www.w3.org/TR/touch-events/

5.2 Participatory Design Sessions

We followed an iterative design process to design our TactualPlot interaction technique for touchscreen chart interaction. We conducted threee participatory design sessions [28] over three weeks, which helped guide our design decisions. The first two sessions lasted 2 hours each, and the final session was 3 hours. The goal was to understand how blind individuals might want to interact with a crossmodal substitution device where tactile input results in auditory output.

Our collaborator is blind, Braille-literate (both reading and writing), and has had recent exposure to tactile graphics and screen readers.

5.3 Dataset and Tasks

We generated 12 datasets to be used in our design sessions and external design reviews. We modeled the data as follows:

- **Data Scatter**: The Pearson's correlation coefficient that varies the strength of the linear relationship between the two variables. We chose three levels for this factor:
  - **Low**: $\rho \in [0.0, 0.30]$
  - **Medium**: $\rho \in (0.3, 0.7]$
  - **High**: $\rho \in (0.7, 1.0]$
- **Data Volume**: The complexity of the dataset is expressed in the number of items being represented in the scatterplots. We chose two levels for this factor:
  - **Small**: 10 items
  - **Medium**: 50 items
- **Polarity of the linear relationship**: This indicates if the linear relationship is positive or negative.

We settled on these values through pilot printing to ensure that we were able to create high-quality tactile points without occlusion. We used the equation: $y = mx + c$, *where* $1 \leq x \leq n$ *and* $n \in \{10, 50\}$ *and* $m \in \{-1, 1, 2, -2\}$, *and* $c \in \{0, 10, 50\}$. Using a negative m value allowed use to generate negative trends in the scatterplot. To introduce a scatter, we added $noise(\varepsilon)$ to both x and y values generated by the linear equation, $where \varepsilon \sim \mathcal{N}(\mu, \sigma^2)$. While changing the slope allows us to control the angle of the trend with reference to the x-axis, we could translate and scatter the points by changing the following values: μ_x, μ_y, σ_x, and σ_y in addition to m and c. We have included our final data specifications and the corresponding (x, y) datasets in the supplementary material. We adapted 3 tasks spanning the 3 categories as described by Sarikaya and Gleicher [46]. We describe the question structure for each of the task types in Table 1.

5.4 Apparatus

We asked our design partner to use both paper-based tactile graphics generated using swell touch paper as well as our prototype implementation of TactualPlot running as a native app using Apache Cordova.

We ran the TactualPlot prototype on an Apple iPad Pro 128 GB with an 11-inch (diagonal) Liquid Retina display (a Liquid Crystal Display, LCD); the actual screen dimensions were 247.6 mm × 178.5 mm. During the test, we ran the native app in fullscreen mode and disabled all notifications.

For both tactile graphs and TactualPlot, we used a square display space measuring approximately 153 × 153 mm. For the tactile graphic, we simply cut each scatterplot to size. For the TactualPlot prototype, we masked the remaining part of the iPad display surface with acrylic plastic cut to size using a laser cutter to prevent accidental touches outside of the scatterplot area. In fact, we placed the iPad underneath the scatterplot even for the tactile graphics condition so that we could track the participant's touch interaction with the tactile graphic as well.

5.5 Design Session 1: Touching Charts

Our goal for the first design session was to understand how chart characteristics such as axis lengths, orientation, data volume would influence data perception. We printed four tactile scatterplots, where

Table 1: **Task types.** List of task types and corresponding question structures for our user study. Each trial corresponded to a given task sub-type.

TASK CATEGORY	TASK TYPE	QUESTION STRUCTURE
Object-centric	Locating (L)	Identify range on Y-axis with maximum value for range between X_1 X_2 on X-axis.
Browsing	Explore trend (T)	As X values increase from left to right, are Y values increasing, decreasing, or random?
Aggregate-level	Numerosity (N)	How many points exist in a specific quadrant: top-left; top-right; bottom-left; or bottom-right?

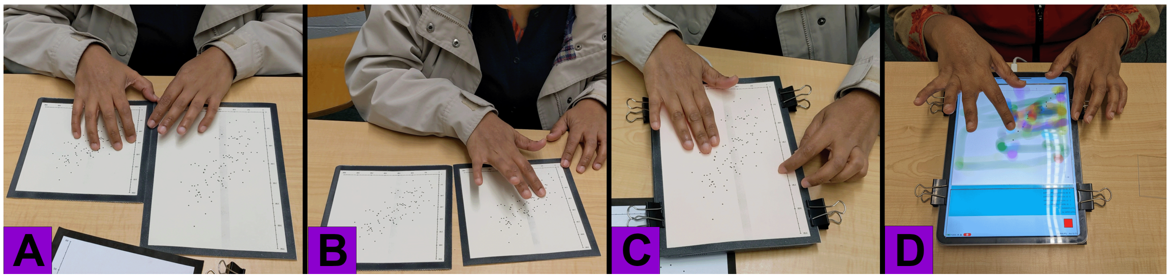

Fig. 3: **Design Sessions 1 and 2.** Our blind collaborator interacting with a low-fidelity tactile graphic using both hands and multiple fingers. (A) Collaborator trying to differentiate between square and rectangular scatterplots that visualize the same dataset in Design Session 1. (B) Collaborator trying to differentiate between two identical scatterplots in Design Session 1. (C) Collecting touch information of our collaborator exploring a tactile scatterplot overlaid on the iPad. (D) Visualizing and observing touch trails of multiple fingers during Design Session 2.

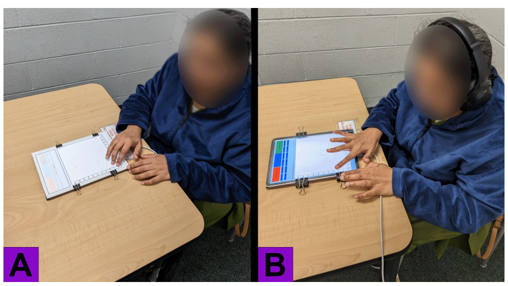

Fig. 4: **Design Session 3.** Our blind collaborator interacting with (A) an iPad-mounted tactile scatterplot, and (B) the TactualPlot prototype.

$N \in \{50, 100\}$ with 2 plots each of uniform axis and non-uniform axis lengths. As seen in Figure 2, we asked our collaborator to identify if the two different datasets were same when printed with uniform vs. non-uniform axis lengths. Just like for visual perception, axis lengths did skew perception. With 100 items and uniform axes, our collaborator sometimes could not tell duplicates apart, based on the exploration route. Overall, we decided use uniform axes to reduce cognitive load.

We had originally designed TactualPlot to support only one-finger interaction. However, it quickly became apparent during the session that our blind collaborator was using not just two hands to interact with the prototypes, but also multiple fingers (Figure 2). They would hold the graphic with one hand, which they reported as a form of physical "anchor" for the other hand. The participant would then use two or even three fingertips of the other hand to slide along the surface of the graphic to sense the embossed data points. This is consistent with work by Wagner et al. [54], which found that people in fact often use not just multiple digits, but sometimes both hands, to interact with tablets. For our touch prototype, such multi-touch interaction was not implemented at the time of testing, and thus did not yield the expected behavior.

This finding also caused us to prioritize such multiple sampling areas rather than speed-dependent dynamic size for a single sampling area. Lastly, we built a test bed to capture touch logs of the exploration route by placing and fastening the tactile graphics to the iPad screen using binder clips. Despite the layer of swell touch paper with raised dots and axes, the iPad touchscreen was able to detect the touches. This motivated us to explore hybrid scatterplots for future versions—a combination of tick mark sonification to indicate that a tick is being touched, and tapping the tick marks to verbalize the tick values.

5.6 Design Session 2: Understanding Multiple Touches

In response to findings from the previous design session, we added multi-touch interaction to the TactualPlot technique so that a user can use several fingertips to spawn multiple sample regions for sensing a larger area on the display. We added brown noise to sonify empty regions by playing a low pitch "buzzing" sound. We experimented with sound envelopes by varying the Attack, Decay, Sustain, and Release (ADSR) values to yield a subjectively pleasing sound for point density.

Our goal was to elicit user exploration strategies [23] given a particular data task. For this session, we used the TactualPlot apparatus (frame **D** in Figure 3), and datasets where $N \in \{10, 50, 100\}$. We first focused on exploring trends (T) in the data using our system. We provided (1) an explanation of the sonification design, (2) verbal description of a linear trend and how that might appear spatially—e.g.,"for a positive trend, imagine that a line is drawn at 45 or 60 degrees with reference to the horizontal axis, and if points are scattered closed to the line, there is a high correlation between x and y", and (3) a recommendation to start with one finger exploration and transition to multiple fingers.

We recorded the iPad screen and implemented a finger trail visualizer that allowed us to study the exploration route for the trends analysis task. Each color represents a a touch using a new finger—a 'new' finger could be the same human finger, i.e., sequential touches with the same index finger is mapped to two different colors. In total, we cycled through 10 different colors to handle 10 concurrent touches (iPad supports 10 touches). Axis, origin, and point density verbalization had not yet been implemented for Session 2.

We include sample exploration strategies to demonstrate the various touch paths that may be possible. In Figure 5, we can see that at the beginning, our collaborator started at the origin and moved horizontally, and moved from the bottom of the canvas to the top. And towards the end of the trial, our participant had sampled all 10 points, but two parts of the canvas remained unexplored (indicated by lack of color). In Figure 6, our collaborator followed a similar trajectory as the previous example, but instead of one finger, they moved two fingers concurrently across the canvas. In this example, all the points were sampled at least once, and only a very small portion of the screen remained unexplored.

In Figure 7, our collaborator placed two fingers at the origin and moved them upwards and then towards the right. After nearing the right extremity of the canvas, our collaborator reestablished contact with the screen at the origin and proceeded to move higher the second time. As seen in the right frame, all 50 points were sampled, but portions of the chart are still unexplored. We considered varying the lens size according to the panning speed, but decided against it. With a fixed

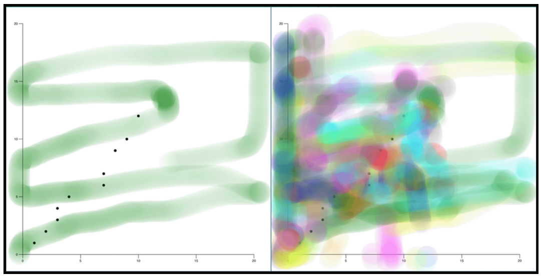

Fig. 5: **Lateral scanning.** *Left:* Initial one finger exploration that starts at the origin, and moves laterally across the screen. *Right:* Final touch trail at the end of the trial.

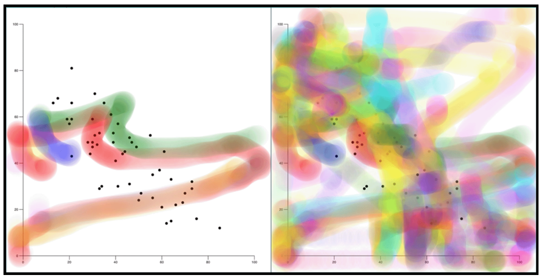

Fig. 6: **Two-finger lateral scanning.** *Left:* Exploration started with two fingers being placed on the origin *Right:* Final touch trail where all the data points have been sampled.

lens size, our participant could detect variations in pitch, but could not accurately map the point density value to a particular pitch value.

5.7 Design Session 3: Exploring Visualization Tasks

We incorporated the feedback from Session 2, and implemented new verbalization features to help better understand "how many points are under your finger" by allowing the user to top on a particular location to verbalize the point density—"1 dot" or "5 dots." Additionally, we included a "clicking" earcon if the user touched a tick mark while panning. We included another discrete touch sound, i.e., a one finger tap on the tick mark would verbalize the tick mark values. For example, "X-axis: 12" or "Y-axis: 10." Our collaborator found the "buzzing" brown noise overwhelming after a while, so we decided to play no sound while continuous dragging on empty space, and instead chose to verbalize "0 dots" on tapping an empty region. We added stereo panning to provide better information of screen location, and played a "click" when the user crossed the mid point of the chart canvas.

Method. For Design Session 3, we wanted to understand if there were any differences in user exploration strategies or perception between TactualPlot and tactile graphics. For a more detailed assessment than the previous two design sessions, we conducted a 3-hour workshop. In addition to the trend task, we asked our collaborator to attempt

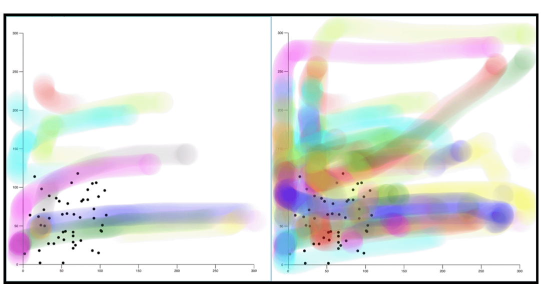

Fig. 7: **Two-finger vertical scanning.** *Left:* Exploration started with two fingers being placed on the origin and moved vertically and then right. *Right:* All the data points have been sampled (3 unexplored regions).

the Numerosity Task (see Table 1). We collected touch logs, qualitative feedback, and task responses. Our collaborated completed the Numerosity tasks using tactile graphics first (see Figure 4A) and then used the TactualPlot apparatus (Figure 4B).

To complete the Numerosity task, our user would have to understand the spatial distribution of the data in addition to perceiving their own fingers' location on screen, and potentially use the axes interactions. To compare both techniques (TactualPlot and tactile graphics) and have time for qualitative feedback, we selected 4 datasets per technique for the numerosity task.

Results. Our collaborator was able to successfully complete all 8 tasks using both tactile graphics and TactualPlot. We report the results in Table 2. We extracted and visualized the touch logs for these tasks in Figure 8. Figure 8 indicates the final state of the system at the end of a trial. We encoded taps (discrete interactions) using a red circle, and continuous dragging in green. A higher opacity of color indicates that a particular region has been sampled more using continuous or discrete interactions. In all 8 trials, the estimated Numerosity value in a particular quadrant was approximately close to the actual value. From the touch logs, we observed that using both techniques, our collaborator was able to restrict finger movement to the quadrant under consideration. Additionally, we observed both from the touch log and qualitative feedback that the sound design for the axes was effective in helping understand finger location. We can observe rough perpendicular finger trails leading to and from the axes, followed or preceded by a discrete tap to hear the tick value. During the session, the researcher observed the interaction style change for the numerosity task; which is reflected in Figure 8.

Observations and Qualitative Feedback. Our collaborator switched from continuous dragging to discrete taps to understand point density because of accidental touches. We removed the "buzzing" sound after the second design session and replaced it with a tap to verbalize "0 dots" for the third session. Continuous dragging, especially using multiple fingers, would often trigger unintended verbalization of the point densities.

The haptic feedback from the tactile dots made it easy to count during continuous dragging, whereas in the TactualPlot system, dragging did not provide exact density estimates: *"no, tactile graphics, I could feel the dots. You know, so that was not my problem, counting them; figuring out—okay, this is, this is the trend, whatever. The audio was a little bit, means, I was not 100% sure. But, it was when you said you know, roughly how many and I said okay, so I counted... one place there's seven, one place was three, then there was two... five... four. That's how I estimated the number."*

When asked if the interface should convey the presence of non-sampled data points—-i.e., points that the user failed to touch—our collaborator indicated that this would help users better understand what they might have missed, especially when analyzing non-synthetic data: *"Oh, you don't need to tell me if I'm looking at real data. Yeah, I'm trying to figure it out, you know, if I pan it like like left and right and whatever... If it has given me the music then I will be able to figure it out that way. Okay, because the points [with] the music in it, it's getting closer and closer."* Our collaborator felt that dataset familiarity may help identify non-sampled data; and explicitly communicating non-sampled data might be helpful to increase confidence. But it's possible that with enough awareness of the data, users may not need the explicit feature to convey the presence of non-sampled points.

6 Design Review

Our goal with TactualPlot is to provide a hybrid representation combining tactile input with sonified output, thereby addressing challenges with both. While we do not anticipate that the approach will be a perfect replacement for a fully dynamic tactile display, we want to provide an acceptable substitute using current touchscreen devices.

To validate the approach, we conducted an expert design review [53] with two blind individuals with domain expertise in statistics and Braille. Our research study was approved both by our university's IRB (ethical review board) as well as the Research Advisory Council of the National Federation of the Blind (NFB) in Baltimore, MD, USA.

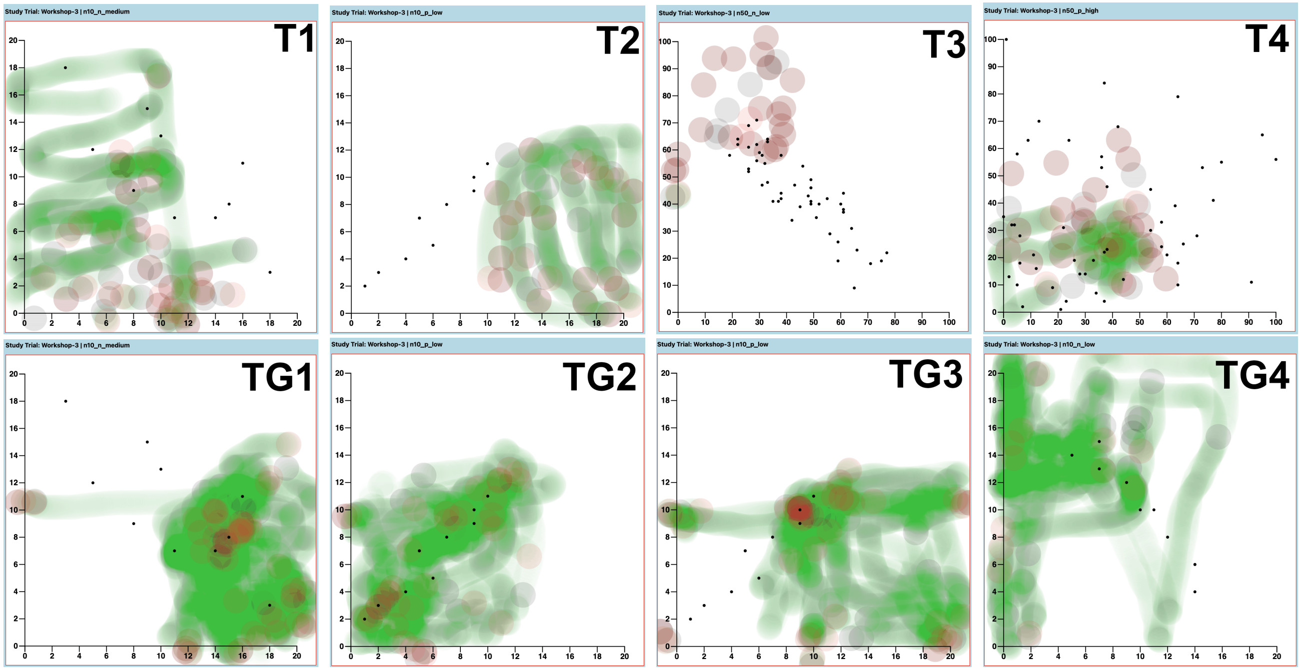

Fig. 8: **Comparison of touch behavior.** Touch interaction for TactualPlot (T) and tactile graphics (TG) for 8 Numerosity (N) task trials. 8 datasets were selected randomly from a set of 12. The red circles represent discrete taps and the green finger trails indicate continuous panning. *First Row (TactualPlot, T1-T4):* Collaborator uses more taps to estimate values in a quadrant, except in T3, where estimates were calculated with only tapping. *Second Row (Tactile Graphics, TG1-TG4):* Collaborator uses more panning to estimate numerosity. We observed that continuous (mapped to non-verbal sound output) and discrete (mapped to speech output) interactions were useful to complete tasks for both plots.

Table 2: **Performance comparison.** Comparing Tactile Graphics and TactualPlot for the Numerosity (N) task.

ID	Data Volume	Data Scatter	Polarity	Quadrant	Estimated Value	Actual Value
Tactile Graphics						
- TG1	10	Medium	N	Bottom-Right	4	5
- TG2	10	Low	P	Bottom-Left	6	9
- TG3	10	Low	P	Top-Right	0	0
- TG4	10	High	N	Top-Left	4	5
TactualPlot						
- T1	10	Medium	N	Bottom-Left	0	0
- T2	10	Low	P	Bottom-Right	1	1
- T3	50	Low	N	Top-Left	20	17
- T4	50	High	P	Bottom-Left	15	24

6.1 Participants

We recruited two participants for our expert design review (2 female) through a local NFB chapter. Demographic details of the reviewers are available in Table 3. Both reviewers were blind screenreader users, and reviewer R1 was also hard of hearing with a Cochlear implant. R1 uses an omni-directional microphone to amplify the sound of the hearing aids. During our session, we connected our system to a Bluetooth speaker and placed it next to her microphone. R1 was recruited for her knowledge of statistics, and R2 for her expertise as a Braille proofreader. The study was conducted in-person because of the need for specialized equipment. To promote participation for people with limited mobility, all sessions were conducted at a public library close to the reviewers' own homes. Reviewers were compensated with a gift card worth $50 for their time and effort.

Table 3: **Expert reviewer demographics and experience.**

ID	Age	Gender	Education	Recruitment criteria
R1	59	Female	Master's degree	Statistics knowledge
R2	60	Female	Bachelor's degree	Braille expertise

6.2 Method and Insights

Tutorial. We introduced our "touching data" presenting tactile graphics and TactualPlot together to allow the reviewers to compare and contrast both approaches. From our formative design sessions, we learned that connecting and abstracting spatial information to scatterplot tasks requires understanding of the visual structure. The researcher explained the concepts of "positive and negative trends", "scatter", and how to perceive a trend with an understanding of point density.

Free Exploration. After receiving confirmation that they had a basic understanding of the sound mapping and interactions, we asked our reviewers to attempt trend (t) tasks 1. Reviewers were asked to freely explore 4 datasets with a low scatter level and both polarities.

We asked our reviewers to think aloud so that the researcher could provide assistance when needed. They could use the equivalent tactile graphic as their "help" system, if needed. At the end of the exploration, the researcher conducted a design debrief to understand how to improve the usability of the TactualPlot technique.

Results. Both our reviewers able to use the features of the system to locate the data points, and perceive some of the semantics of the scatterplot such as angles from the reference axes, clustering and spread along a linear line: (R1) *"Ah, yes. So once you find the one dot, you try to basically hold... You have one hand... hold the dot, and then go over to the Y-axis and then one down to the X-axis. Just as, as you normally would with a protractor back in the day."* With more time spent on free exploration, R1 started abstracting the spatial and sound information towards identifying the data trend: *"It [pitch]] doesn't seem to be going up as fast, and it [pitch] starts out lower. The sound is not as as high. At least to me, to me. So, I don't think it seems. And because it sounds lower, I think that the plot is more spread out."*

R2 highlighted the trade offs between tactile plots and TactualPlot and believed that she prefers tactile plots because she is more attuned to "reading quickly" with multiple fingers. R2 pointed out the difficulty in understanding which finger was producing the sound when many fingers are touching the screen.

Both reviewers needed more time to effectively interpret crossmodal substitution: *"So, but I think intuitively, it, it has a lot of sense. If, you know, you can just, you know, I think, get used to some, some of maybe the technical things, like I, I'm, you know, I didn't get a computer until I was like, in my 30s. So, some of it, it may be just more like technical."*

7 Discussion

Our longitudinal participatory design method allowed us to make better design decisions towards discovering the strengths and limitations of crossmodal substitutions. Secondly, our comparison to tactile graphics provided an opportunity to better understand touch-bases data exploration strategies which are crucial in helping with graphical perception in touchscreens [23]. Below, we describe the challenges that our participants faced while using our technique and discuss solutions that can be introduced to improve touchscreen data accessibility.

7.1 Improving Crossmodal Sensory Substitution

Continuously touching data could be better for understanding trends, but for certain tasks where accuracy of the touch location matters, we noticed accidental touches that could add complexity to interpreting the data. Discrete actions on the other hand are slower but provide more control. Accidental touches can be a problem in sonified scatter plots, as they can lead to unintended disruptions during data exploration. Additionally, discrete interactions could be leveraged towards providing more granularity [61] by allowing the user to drill-down into the data. Introducing a grid structure [7] and other scaffolds such as rulers can help users understand the spatial relationships between data points more easily. Guiding sounds can be played to ensure that users confidently sample all the data points necessary to complete a data task [40]. For example, you could imagine designing an audio version of the Halo [6] technique that uses concentric rings to display off-screen targets.

Tactile graphics can provide better parallel access to data, allowing blind users to explore the plot more efficiently. As we observed during our assessments, careful sound design is needed to ensure that users are not overwhelmed by sound and speech. Designing pleasant and engaging soundscapes can help improve the user experience.

Voice interactions and question-answering [32] can be a useful addition to our technique; allowing users to interact with the plot using spoken commands or queries for data related tasks while relying on touch to perceive the chart structure. For large-scale data sets, it may be useful to provide aggregate functions, such as summary statistics or clustering, to help users identify patterns and trends in the data. In addition to providing more information, listening to summary statistics may also improve confidence during data exploration.

Our results show that understanding data spread is achievable by listening to sonified data densities. Initially, we mapped individual point values to pitches, and realized that simultaneously playing multiple pitches (tones) on many data points being touched led to cacophony. We use non-verbal sound is used to convey spatial patterns and help identify data trends, while data values and counts are verbalized through speech. Data values can be mapped to acoustic parameters such as frequency, loudness, timbre, or tempo; but sonification design needs to be empirically evaluated for blind individuals owing to varying sound perception abilities [55, 56]. We recommend empirically designing and evaluating sound design while employing audio parameters, such as spatial audio for localization, familiar sounds [27], multiple timbres and acoustic dimensions, to improve visualization task performance [58] for chart representations that use multiple visual channels.

Finally, another low-cost and readily available modality that we could explore—either in in conjunction with or as a replacement for sound—is vibrotactile feedback [12, 42]. Vibration may serve as an even better output channel because touching a high-density area of data would seemingly yield a large amount of "friction." Specialized vibrotactile hardware could go even further; for example, the HoloBraille system [38] uses a custom-made case to independently convey vibrotactile output for six areas representing the dots in a Braille character.

7.2 Limitations and Future Work

There are several caveats and limitations with our work. For example, the number of participants in our evaluation was small (one in our formative design, and two in our expert review). While this is a fair point, our recent experiences in designing assistive technology for accessible visualization tell us to move slowly and deliberately, ensuring buy-in from the blind community at all times. In retrospect, the insights gained in this research are a direct result of the careful participatory design method. While we are confident in the state of TactualPlot, we also feel that the technique would benefit from a wider and more quantitative evaluation involving a diverse set of participants. In fact, given the software-only nature of the tool, we believe a large-scale crowdsourced study would eventually be possible once we improve the interface, and also develop the right scaffolding to teach spatial data perception using crossmodal substitution. Additionally, the stimuli that we used did not contain real data, and the sound design was not experimentally validated. In the future, we plan to experimentally determine how spatial audio can be used to help users navigate and interpret chart data.

In this paper, we only collected data that helped us understand low-level data tasks. Future work can focus on exploring touchscreen and other system logs to understand how blind users explore the chart canvas during both low-level visualization and high-level visualization tasks [8]. Furthermore, we hope to tackle statistical graphics such as barcharts, piecharts, and linecharts in the future. For all of these charts, we would utilize the same crossmodal approach where the user's fingertips are used to "touch" data points represented using these different graphics, as if they were tactile graphics.

Another potential weakness is that despite our best efforts, exploration using tactile graphics made our participants more confident in their answers when compared to TactualPlot. However, this is not surprising; paper-based tactile graphics are essentially best in class for perceiving spatial data representations, and they have an important limitation: they are not dynamic. For comparison, we ended up spending approximately $30 ($1.50 per sheet) in total on the swell touch paper used in this study, and printing a single sheet was both cumbersome and inconvenient. Once refreshable tactile displays such as the DTD are widely available, the need for software-only tools such as TactualPlot might diminish. However, that day is still in the future, and even when it comes, their cost will likely exceed existing smartphone devices.

8 Conclusion

We have presented TactualPlot, a multitouch-actuated scatterplot display for blind users that yields data density sonified using pitch. This crossing from one sense—touch—to another—sound—is an example of crossmodal sensory substitution. The primary contribution of our work is the in-depth design process and qualitative validation that we conducted for TactualPlot. Together with our blind collaborator, we have described this process of designing a novel form of assistive technology for data visualization. We derive a technique that can be used to design touch-based interactions for 2D scatterplots on touchscreen devices. We feel that our account in this paper contributes to the science and design of accessible data visualization, and we are eager to see how our findings impact the community.

Acknowledgments

This material is based upon work supported by the U.S. National Science Foundation under Grant No. IIS-2211628. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency. We thank Eric Newburger for guidance on generating datasets, and Biswaksen Patnaik for assistance with laser cutting.

References

[1] C. H. Abraham, B. Boadi-Kusi, E. K. A. Morny, and P. Agyekum. Smartphone usage among people living with severe visual impairment and blindness. *Assistive Technology*, 34(5):611–618, 2022. doi: 10.1080/10400435.2021.1907485 1, 3

[2] R. A. Al-Zaidy and C. L. Giles. Automatic extraction of data from bar charts. In *Proceedings of the ACM Conference on Knowledge Capture*, pp. 30:1–30:4. ACM, 2015. doi: 10.1145/2815833.2816956 2

[3] J. Alexander, A. Roudaut, J. Steimle, K. Hornbæk, M. B. Alonso, S. Follmer, and T. Merritt. Grand challenges in shape-changing interface research. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 299:1–299:14. ACM, New York, NY, USA, 2018. doi: 10.1145/3173574.3173873 1, 3

[4] P. Bach-y-Rita and S. W. Kercel. Sensory substitution and the human–machine interface. *Trends in Cognitive Sciences*, 7(12):541–546, 2003. doi: 10.1016/j.tics.2003.10.013 1, 2, 3

[5] A. Batch, B. Patnaik, M. Akazue, and N. Elmqvist. Scents and sensibility: Evaluating information olfactation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 1–14. ACM, New York, NY, USA, 2020. doi: 10.1145/3313831.3376733 2

[6] P. Baudisch and R. Rosenholtz. Halo: a technique for visualizing off-screen objects. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 481–488. ACM, New York, NY, USA, 2003. doi: 10.1145/642611.642695 9

[7] M. Blanco, J. Zong, and A. Satyanarayan. Olli: An Extensible Visualization Library for Screen Reader Accessibility. In *IEEE VIS Posters*, 2022. 9

[8] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, 2013. doi: 10.1109/TVCG.2013.124 9

[9] S. Cantrell, B. N. Walker, and Ø. Moseng. Highcharts sonification studio: An online, open-source, extensible, and accessible data sonification tool. In *Proceedings of the International Conference on Auditory Display*, 2021. 2

[10] T. Carter, S. A. Seah, B. Long, B. W. Drinkwater, and S. Subramanian. UltraHaptics: multi-point mid-air haptic feedback for touch surfaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 505–514. ACM, New York, NY, USA, 2013. doi: 10.1145/2501988.2502018 3

[11] J. Choi, S. Jung, D. G. Park, J. Choo, and N. Elmqvist. Visualizing for the non-visual: Enabling the visually impaired to use visualization. *Computer Graphics Forum*, 38(3):249–260, 2019. doi: 10.1111/cgf.13686 2

[12] S. Choi and K. J. Kuchenbecker. Vibrotactile display: Perception, technology, and applications. *Proceedings of the IEEE*, 101(9):2093–2104, 2013. doi: 10.1109/JPROC.2012.2221071 9

[13] P. Chundury, B. Patnaik, Y. Reyazuddin, C. Tang, J. Lazar, and N. Elmqvist. Towards understanding sensory substitution for accessible visualization: An interview study. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1084–1094, 2022. doi: 10.1109/TVCG.2021.3114829 1, 2, 4

[14] A. M. Cook and J. M. Polgar. *Assistive Technologies: Principles and Practice*. Elsevier Health Sciences, 4th ed., 2014. 1, 2, 3

[15] O. Deroy and M. Auvray. Reading the world through the skin and ears: A new perspective on sensory substitution. *Frontiers in Psychology*, 3, 2012. doi: 10.3389/fpsyg.2012.00457 2, 3

[16] A. Drogemuller, A. Cunningham, J. A. Walsh, J. Baumeister, R. T. Smith, and B. H. Thomas. Haptic and visual comprehension of a 2D graph layout through physicalisation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 463:1–463:16. ACM, New York, NY, USA, 2021. doi: 10.1145/3411764.3445704 1

[17] F. Elavsky, C. L. Bennett, and D. Moritz. How accessible is my visualization? evaluating visualization accessibility with Chartability. *Computer Graphics Forum*, 41(3):57–70, 2022. doi: 10.1111/cgf.14522 2

[18] N. Elmqvist. Visualization for the blind. *ACM Interactions*, 30(1):52–56, 2023. doi: 10.1145/3571737 2, 3

[19] K. Enge, A. Rind, M. Iber, R. Höldrich, and W. Aigner. Towards multimodal exploratory data analysis: SoniScope as a prototypical implementation. In *Short Paper Proceedings of the IEEE VGTC/Eurographics Conference on Visualization*, pp. 67–71. Eurographics Association, 2022. doi: 10.2312/evs.20221095 4

[20] D. Fan, A. F. Siu, W. A. Law, R. R. Zhen, S. O'Modhrain, and S. Follmer. Slide-tone and tilt-tone: 1-DOF haptic techniques for conveying shape characteristics of graphs to blind users. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2022. doi: 10.1145/3491102.3517790 2

[21] W. W. Gaver, A. Dunne, and E. Pacenti. Design: Cultural probes. *ACM Interactions*, 6(1):21–29, 1999. doi: 10.1145/291224.291235 5

[22] C. Gleason, P. Carrington, C. T. Cassidy, M. R. Morris, K. M. Kitani, and J. P. Bigham. "It's almost like they're trying to hide it": How user-provided image descriptions have failed to make Twitter accessible. In *Proceedings of the World Wide Web Conference*, pp. 549–559. ACM, New York, NY, USA, 2019. doi: 10.1145/3308558.3313605 2

[23] J. L. Gorlewicz, J. L. Tennison, P. M. Uesbeck, M. E. Richard, H. P. Palani, A. Stefik, D. W. Smith, and N. A. Giudice. Design guidelines and recommendations for multimodal, touchscreen-based graphics. *ACM Transactions on Accessibile Computing*, 13(3):10:1–10:30, 2020. doi: 10.1145/3403933 6, 9

[24] D. Guinness, A. Muehlbradt, D. Szafir, and S. K. Kane. RoboGraphics: Dynamic tactile graphics powered by mobile robots. In *Proceedings of the ACM Conference on Computers and Accessibility*, pp. 318–328. ACM, New York, NY, USA, 2019. doi: 10.1145/3308561.3353804 3

[25] T. Hermann, A. Hunt, and J. G. Neuhoff. *The Sonification Handbook*. Logos Verlag, Berlin, Germany, 2011. 2

[26] L. Holloway, K. Marriott, and M. Butler. Accessible maps for the blind: Comparing 3D printed models with tactile graphics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 198:1–198:13. ACM, New York, NY, USA, 2018. doi: 10.1145/3173574.3173772 2, 3

[27] M. N. Hoque, M. Ehtesham-Ul-Haque, N. Elmqvist, and S. M. Billah. Accessible data representation with natural sound. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 826:1–826:19. ACM, New York, NY, USA, 2023. doi: 10.1145/3544548.3581087 2, 4, 9

[28] S. Jänicke, P. Kaur, P. Kuzmicki, and J. Schmidt. Participatory visualization design as an approach to minimize the gap between research and application. In *Proceedings of the Eurographics Workshop on the Gap between Visualization Research and Visualization Software*, pp. 35–42. Eurographics Association, 2020. doi: 10.2312/visgap.20201108 5

[29] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, and K. Hornbæk. Opportunities and challenges for data physicalization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 3227–3236. ACM, New York, NY, USA, 2015. doi: 10.1145/2702123.2702180 1, 3

[30] C. Jung, S. Mehta, A. Kulkarni, Y. Zhao, and Y. Kim. Communicating visualizations without visuals: Investigation of visualization alternative text for people with visual impairments. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1095–1105, 2022. doi: 10.1109/TVCG.2021.3114846 2

[31] G. Kim, J. Kim, and Y.-S. Kim. Explain what a treemap is: Exploratory investigation of strategies for explaining unfamiliar chart to blind and low vision users. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 805:1–805:13. ACM, New York, NY, USA, 2023. doi: 10.1145/3544548.3581139 2

[32] J. Kim, A. Srinivasan, N. W. Kim, and Y. S. Kim. Exploring chart question answering for blind and low vision users. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 828:1–828:15. ACM, New York, NY, USA, 2023. doi: 10.1145/3544548.3581532 9

[33] G. Kramer, ed. *Auditory Display: Sonification, Audification, and Auditory Interfaces*, vol. XVIII of *Santa Fe Institute Studies in the Sciences of Complexity*. Addison-Wesley, Reading, MA, USA, 1994. 2

[34] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. H. Flowers, N. Miner, J. Neuhoff, R. Bargar, S. Barrass, J. Berger, G. Evreinov, W. T. Fitch, M. Gröhn, S. Handel, H. Kaper, H. Levkowitz, S. Lodha, B. Shinn-Cunningham, M. Simoni, and S. Tipei. Sonification report: Status of the field and research agenda. Technical report, National Science Foundation, 1997. 2

[35] A. Lundgard, C. Lee, and A. Satyanarayan. Sociotechnical considerations for accessible visualization design. In *Proceedings of the IEEE Visualization Conference*, pp. 16–20. IEEE Computer Society, Los Alamitos, CA, USA, 2019. doi: 10.1109/VISUAL.2019.8933762 1, 2

[36] K. Nakagaki, L. Vink, J. Counts, D. Windham, D. Leithinger, S. Follmer, and H. Ishii. Materiable: Rendering dynamic material properties in response to direct physical touch with shape changing interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 2764–2772. ACM, New York, NY, USA, 2016. doi: 10.1145/2858036.2858104 1

[37] B. Nanay. Sensory substitution and multimodal mental imagery. *Perception*, 46(9):1014–1026, 2017. doi: 10.1177/0301006617699225 2, 3

[38] H. Nicolau, K. Montague, T. J. V. Guerreiro, A. Rodrigues, and V. L. Hanson. HoliBraille: multipoint vibrotactile feedback on mobile devices. In *Proceedings of the ACM Web for All Conference*, pp. 30:1–30:4. ACM, New York, NY, USA, 2015. doi: 10.1145/2745555.2746643 9

[39] D. Nikitenko and D. J. Gillis. Touching the data: Exploring data sonification on mobile touchscreen devices. In *Proceedings of the International Conference on Mobile Systems and Pervasive Computing*, vol. 34 of *Procedia Computer Science*, pp. 360–367. Elsevier, 2014. doi: 10.1016/j.procs.2014.07.038 3

[40] U. Oh, S. K. Kane, and L. Findlater. Follow that sound: using sonification and corrective verbal feedback to teach touchscreen gestures. In *Proceedings of the ACM Conference on Computers and Accessibility*, pp. 13:1–13:8. ACM, 2013. doi: 10.1145/2513383.2513455 9

[41] D. Osiński, M. Łukowska, D. R. Hjelme, and M. Wierzchoń. Colorophone 2.0: A wearable color sonification device generating live stereo-soundscapes—design, implementation, and usability audit. *Sensors*, 21(21), 2021. doi: 10.3390/s21217351 2

[42] H. P. Palani, P. D. S. Fink, and N. A. Giudice. Design guidelines for schematizing and rendering haptically perceivable graphical elements on touchscreen devices. *International Journal of Human–Computer Interaction*, 36(15):1393–1414, 2020. doi: 10.1080/10447318.2020.1752464 9

[43] B. Patnaik, A. Batch, and N. Elmqvist. Information olfactation: Harnessing scent to convey data. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):726–736, 2019. doi: 10.1109/TVCG.2018.2865237 2

[44] V. Potluri, J. Thompson, J. Devine, B. Lee, N. Morsi, P. de Halleux, S. Hodges, and J. Mankoff. PSST: enabling blind or visually impaired developers to author sonifications of streaming sensor data. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 46:1–46:13. ACM, New York, NY, USA, 2022. doi: 10.1145/3526113.3545700 2

[45] K. Salisbury, F. Conti, and F. Barbagli. Haptic rendering: Introductory concepts. *IEEE Computer Graphics and Applications*, 24(2):24–32, 2004. doi: 10.1109/MCG.2004.1274058 1, 3

[46] A. Sarikaya and M. Gleicher. Scatterplots: Tasks, data, and designs. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):402–412, 2018. doi: 10.1109/TVCG.2017.2744184 5

[47] A. Sharif, O. H. Wang, and A. T. Muongchan. "what makes sonification user-friendly?" exploring usability and user-friendliness of sonified responses. In *Proceedings of the ACM Conference on Computers and Accessibility*, pp. 45:1–45:5. ACM, New York, NY, USA, 2022. doi: 10.1145/3517428.3550360 4

[48] A. Sharif, O. H. Wang, A. T. Muongchan, K. Reinecke, and J. O. Wobbrock. VoxLens: Making online data visualizations accessible with an interactive JavaScript plug-in. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 478:1–478:19. ACM, New York, NY, USA, 2022. doi: 10.1145/3491102.3517431 2

[49] B. Shneiderman. Direct manipulation: A step beyond programming languages. *Computer*, 16(8):57–69, 1983. doi: 10.1109/MC.1983.1654471 3

[50] A. F. Siu, G. S. Kim, S. O'Modhrain, and S. Follmer. Supporting accessible data visualization through audio data narratives. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 476:1–476:19. ACM, 2022. doi: 10.1145/3491102.3517678 2

[51] F. Taher, Y. Jansen, J. Woodruff, J. Hardy, K. Hornbæk, and J. Alexander. Investigating the use of a dynamic physical bar chart for data exploration and presentation. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):451–460, 2017. doi: 10.1109/TVCG.2016.2598498 1, 3

[52] J. R. Thompson, J. J. Martinez, A. Sarikaya, E. Cutrell, and B. Lee. Chart Reader: Accessible visualization experiences designed with screen reader users. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 802:1–802:18. ACM, New York, NY, USA, 2023. doi: 10.1145/3544548.3581186 2

[53] M. Tory and T. Möller. Evaluating visualizations: Do expert reviews work? *IEEE Computer Graphics and Applications*, 25(5):8–11, 2005. doi: 10.1109/MCG.2005.102 7

[54] J. Wagner, S. Huot, and W. E. Mackay. BiTouch and BiPad: designing bimanual interaction for hand-held tablets. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 2317–2326. ACM, New York, NY, USA, 2012. doi: 10.1145/2207676.2208391 6

[55] B. N. Walker and G. Kramer. Mappings and metaphors in auditory displays: An experimental assessment. *ACM Trans. Appl. Percept.*, 2(4):407–412, 2005. doi: 10.1145/1101530.1101534 9

[56] B. N. Walker and L. M. Mauney. Universal design of auditory graphs: A comparison of sonification mappings for visually impaired and sighted listeners. *ACM Transactions on Accessibile Computing*, 2(3), 2010. doi: 10.1145/1714458.1714459 9

[57] J. Wallace, J. C. McCarthy, P. C. Wright, and P. Olivier. Making design probes work. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 3441–3450. ACM, New York, NY, USA, 2013. doi: 10.1145/2470654.2466473 5

[58] R. Wang, C. Jung, and Y. Kim. Seeing through sounds: Mapping auditory dimensions to data and charts for people with visual impairments. *Computer Graphics Forum*, 41(3):71–83, 2022. doi: 10.1111/cgf.14523 2, 9

[59] D. Yau and R. Mccrindle. Musicam: an instrument to demonstrate chromaphonic synesthesia. *Digital Creativity*, 18(2):121–127, 2007. doi: 10.1080/14626260701401510 2

[60] H. Zhao, C. Plaisant, B. Shneiderman, and J. Lazar. Data sonification for users with visual impairment: A case study with georeferenced data. *ACM Transactions on Computer-Human Interaction*, 15(1):4:1–4:28, 2008. doi: 10.1145/1352782.1352786 2

[61] J. Zong, C. Lee, A. Lundgard, J. Jang, D. Hajas, and A. Satyanarayan. Rich screen reader experiences for accessible data visualization. *Computer Graphics Forum*, 41(3):15–27, 2022. doi: 10.1111/cgf.14519 2, 9