Design Guidelines and Recommendations for Multimodal, Touchscreen-based Graphics

JENNA L. GORLEWICZ and JENNIFER L. TENNISON, Saint Louis University
P. MERLIN UESBECK, University of Nevada, Las Vegas
MARGARET E. RICHARD, Saint Louis University
HARI P. PALANI, VEMI Lab, University of Maine
ANDREAS STEFIK, University of Nevada, Las Vegas
DERRICK W. SMITH, University of Alabama in Huntsville
NICHOLAS A. GIUDICE, VEMI Lab, University of Maine

With content rapidly moving to the electronic space, access to graphics for individuals with visual impairments is a growing concern. Recent research has demonstrated the potential for representing basic graphical content on touchscreens using vibrations and sounds, yet few guidelines or processes exist to guide the design of multimodal, touchscreen-based graphics. In this work, we seek to address this gap by synergizing our collective research efforts over the past eight years and implementing our findings into a compilation of recommendations, which we validate through an iterative design process and user study. We start by reviewing previous work and then collate findings into a set of design guidelines for generating basic elements of touchscreen-based multimodal graphics. We then use these guidelines to generate exemplary graphics in mathematics, specifically bar charts and geometry concepts. We discuss the iterative design process of moving from guidelines to actual graphics and highlight challenges. We then present a formal user study with 22 participants with visual impairments, comparing learning performance on using touchscreen-rendered graphics to embossed graphics. We conclude with qualitative feedback from participants on the touchscreen-based approach and offer areas of future investigation as these recommendation are expanded to include more complex graphical concepts.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**; **User studies**; **Touch screens**; **Haptic devices**; *Usability testing*; *Accessibility technologies*; *Accessibility systems and tools*; • **Social and professional topics** → People with disabilities;

1 MOTIVATION

Access to graphics remains one of the most pervasive information access challenges faced by in-
dividuals with visual impairments (VI) today. Lack of access to graphical information has implica-
tions in social, vocational, and educational prospects. It is estimated that only 11% of individuals
with VI hold bachelor's degrees [19] and that over 60% of working-age individuals with VI are
unemployed or under-employed [3].

Over 12M people in the U.S. and 285M people worldwide are estimated as having some form of
significant and uncorrected visual impairment. This number is expected to grow as our population
ages [82]. Unfortunately, the gap in information access between sighted and individuals with VI is
widening. Because it is often less expensive and easier to deploy content digitally, there has been a
strong movement away from text-based materials to electronic, graphic-rich content. Increasingly,
highly visual content such as graphs, diagrams, charts, images, maps, and photographs are being
displayed using visual-based digital technologies. Consider photographs on social media feeds,
the images and advertisements on web pages, or the figures in e-books. While this content can
be represented via nonvisual means, such as a text-based description, these nonvisual solutions
often do not capture the rich information conveyed in the graphic and provide little insight on the
structure of how graphical information is spatially organized. Further, books such as Reference [1]
that have multimodal, digital accessibility features, are rare.

For individuals with VI, tactile representations are often critical to garnering adequate spatial
information. However, meaningful tactile representations of graphics often require the use of an
embosser (e.g., View Plus [57]), swell paper [75], or other creative solutions that involve tactile ma-
nipulatives [12, 58]. While these solutions are readily available and widely used, they often require
specialized personnel and hardware for authoring and producing content. The output created is
also static renderings, meaning that a new image must be created every time there is a change
or update to the graphic. Tactile embossers in particular are relatively expensive (normally rang-
ing from $5000–$15,000), causing adoption to be contained largely within educational or support
settings as opposed to household and common consumer settings [17, 57].

The evolution of dynamic touch-based interfaces has alleviated some of these barriers. For in-
stance, a host of refreshable tactual technologies have been developed based on a variety of dif-
ferent schemes: force feedback [21, 47, 69, 77, 84, 86, 87], refreshable pin arrays [4, 70], micro flu-
idics [64], moldable alloys, and surface haptics [21, 47, 69, 77, 84], to name a few (see Reference [51]
for a review). Investigations have also explored repurposing existing hardware, such as the pen
plotter in Reference [46], to create accessible graphics through audio and kinesthetic feedback.
While such developments are the driving force behind creating new haptic technology to provide
better access, these solutions are not widely available nor broadly adopted. This is likely due to
several factors including the high cost and lack of commercial availability associated with most of
the haptic systems, the in-depth manufacturing and fabrication process required for some of the
technologies, and the need for additional hardware that only adds to the host of access devices and
technologies already used by individuals with VI.

2 BACKGROUND AND RELATED WORK

Our group has investigated the use of vibrations and sounds generated by commercially available touchscreens as another solution to help bridge the graphical information gap. Touchscreen-based smart devices are increasingly becoming the de facto platform for accessing and interacting with digital information due to their convenience and widespread availability [56], which is fueling improved accessibility functions. While touchscreens are still largely used visually, advances in the constituent software and hardware have begun to increase the accessibility of such displays for nonvisual information access [5, 65]. Built-in screen reading software (e.g., Apple VoiceOver and Google TalkBack) leveraging text-to-speech engines have provided auditory access to text-based material. Additional accessibility options embedded within the native operating systems enable users to interact with on-screen content via speech, gestures, or through contrast and color-change options. In addition to these built-in accessibility advancements, early research on making touchscreen platforms more accessible has shown great promise. Auditory-based solutions, such as that in References [6, 7, 37, 50, 67], demonstrate the enhancements that auditory feedback alone can provide in accessing touchscreen applications or teaching gestures. Software-based access overlays have also been put forth to increase nonvisual access to on-screen content [38], particularly for larger touch surfaces. All of these efforts support the potential of touchscreens to be accessible technologies.

Advancements in touchscreen hardware for accessibility are slower to fruition. While promising touchscreen-based technologies such as microfluidic displays [51, 64] and surface haptic displays [21, 47, 69, 77, 84] are being developed, many of these are still in the research phase and suffer from high component costs or reliance on hardware-specific platforms. Yet, there may be a more streamlined solution already available, as touchscreens are now often outfitted with vibration capabilities. Though vibrations often lack inherent, intuitive physical properties, research has demonstrated that vibration feedback promotes notable enhancements to the touchscreen user experience [14, 34]. Early investigations into touchscreen vibrations have demonstrated that users can type more accurately and can complete tasks such as dragging and dropping or selecting objects on-screen more efficiently [11, 32]. User experience has demonstrated the value of vibrations in alerting and cuing—whether it be for signaling the arrival of incoming content or user-specified notifications [34]. While this secondary form of feedback has improved the overall touchscreen experience, it has done little to improve robust accessibility of information content, specifically graphics-based content. Indeed, vibrations have often been relegated to tertiary use on touchscreens. This is in spite of research demonstrating that vibrations can represent semantic information [35, 85] and even tactile language building blocks [10]. For more information on vibrotactile perception and technology advancements, we refer readers to References [14, 40, 44].

While it is acknowledged that there are limits to the information conveyed using vibrations, our group has hypothesized that more can be done with vibrations as a primary means of communicating on-screen information. This hypothesis is supported by recent research demonstrating the use of vibrations for supporting navigation, education, and everyday tasks for individuals with VI [23–25, 39]. Particularly in education and spatial learning, recent investigations have demonstrated that vibrations are effective in representing maps [59] and building floor plans [22], accurately perceiving and following lines and graphs [23, 55, 73], representing graphical shapes [24, 72], and representing simple graphic entities [30]. Also, vibrations have been shown to aid in non-visual panning and zooming of information on the screen [52, 54]. Additional discussions on the use of vibrations for supporting accurate non-visual graphic rendering on touchscreens can be found in References [26, 30, 39]. These findings have illustrated the potential for vibrations to not only be a primary interaction modality for all individuals, but to be a critical piece in closing the graphical

access gap for individuals with VI. Namely, because vibration actuation is common in commercially available, touch-based platforms, the necessary hardware capabilities are already present and adopted within the VI community, making the potential for adoption and widespread use high. This approach, albeit not as high in resolution due to small screen sizes and not as haptically salient as traditional, embossed graphics, has many benefits, with one of the primary advantages being that it can almost immediately reach the hands of a high percentage of end-users. Recent estimates suggest that 70%–90% of VI cell phone users have smart phones in the United States. This high penetration of the VI AT market is incredibly challenging to attain with single-purpose, hardware solutions. Moreover, these platforms are inherently multisensory, using visual, auditory, and tactile feedback to convey information, which better serves all users.

Despite this leap forward in using haptic feedback to increase the accessibility of information, the proposed touchscreen solution does not come without limitations. Commercially available touchscreens are featureless glass panes currently incapable of delivering localized vibration. This is problematic, because it provides users with no gradient of direction and can trigger sensory adaptation of feedback, fatiguing a user's finger of the vibration sensation [26, 39, 55]. Whereas a user is able to use subtle tactile cues on embossed graphics, such as feeling that a line segment is on the left side of the fingertip, vibratory touchscreens do not offer such fine distinction and, instead, vibrate the entirety of the fingertip no matter what part is in contact with an object. This saturation of sensation can then lead to sensory fatigue and adaptation, making it difficult to continue to sense vibrating features on the touchscreen. Touch, unlike vision, also lacks specialized processors for parallel information and the user must undergo a complicated process of synthesizing touch information. Users must keep track of all important kinesthetic and spatial cues to perceive content correctly, highlighting the haptic system's spatial bandwidth and temporal processing limitations [39].

Additionally, creating and authoring multisensory content is a new area of research. While empirical parameters exist for visual rendering of graphics, few basic parameters are available for rendering multimodal graphics, particularly those that are more complex, on touchscreens. Few guidelines exist to aid designers in applying vibration and sound feedback to an image for optimal consumption. This article serves to address this gap and to disseminate the intricacies of designing multimodal images that may not always be detailed in published research. First, we collate and compile findings of our own and other researchers' work in this area, formulating design guidelines for representing fundamental graphical elements multimodally on touchscreens. Second, we present two case studies of applying these guidelines to more complex mathematical graphics to illustrate the iterative design process behind creating multimodal graphics. Third, we provide a formal user study with 22 participants with VI, comparing user performance on touchscreen-rendered graphics created using the above guidelines and process with embossed graphics created using tactile graphic standards. We conclude by sharing insights for designers and practitioners in creating multisensory graphics and offer suggestions for future areas of investigation to propel this idea into broader adoption.

3   DESIGN GUIDELINES AND RECOMMENDATIONS

While touch is a fundamental part of our perception [42, 45], it is not an all-encompassing substitute for vision. Haptic perception differs from vision in important ways, as vision is estimated to have 500 times greater sensory bandwidth than touch [43]. Differences between these modalities occur in aspects such as how the information is accessed, the type and amount of access to data that is available, sensory adaptation, spatial resolving power, temporal integration, spatial localization, and vulnerability to systematic distortion [31, 39]. These limitations undoubtedly challenge designers of multimodal, tactile graphics.

Unlike traditional tactile graphics, which are composed of physical deformations such as raised dots, touchscreen-based graphics are conveyed via vibrations and auditory/speech cues on a flat surface. A primary challenge in designing multimodal, vibrotactile-based graphics is that we lack guidelines that inform perception and usability of these graphics as displayed on screen. For example, what size should basic graphical elements (points, lines) be? What elements should be represented with what type of feedback? Answering such questions is critical to informing the specification of these needed guidelines. Further, while there is significant literature on tactile thresholds for physical stimuli [20, 36], the effect of aging on tactile ability [66], and the tactile "science" supporting Braille reading and pattern recognition [76], vibrotactile perception differs in several ways compared to the stimuli received from physical tactile graphics. Namely, physical stimuli are often cutaneous cues perceived via pressure-based mechanoreceptors, whereas vibration cues innervate pacinian corpuscles, which are functionally different. Because of this distinction, findings from previous research using physical stimuli are not necessarily relevant for rendering touchscreen-based vibrotactile graphics.

While there are principles that exist for creating hard-copy tactile graphics [8], only a subset of them apply here, because of the fundamental differences in the way these graphics are conveyed. Nonetheless, guidelines and standards for tactile graphics such as those developed by the Braille Authority of North America [8] and Braille Literacy Canada [9] have formalized hard-copy tactile graphics and have provided a framework for consistency in presentation of information that can be widely adopted. A similar framework is needed for vibrotactile, touchscreen-based graphics, such that these graphics can be scoped appropriately, represented consistently, and interpreted efficiently. To address the above challenge, our group has collectively probed several perceptual and usability factors that inform how basic graphical components should be represented on vibrotactile interfaces, with a particular focus on STEM-based (Science, Technology, Engineering, and Math) graphics. This article collates our findings across 8+ years of study to put forth recommendations, design guidelines, and best practices for rendering graphics multimodally on touchscreen platforms.

In the sections that follow, we will discuss our recommendations for the creation and usability of multimodal, touchscreen-based graphics. The sections will follow the outline presented in Figure 1 and discuss the representation of fundamental graphical elements, how to assign feedback, suggestions for hardware modification, and what common user strategies to account for when designing material.

3.1 Fundamental Graphical Elements

Points, lines, and polygons represent some of the most basic graphical elements from which more complex graphics (such as charts, diagrams, and graphs) are created. As such, our group's early investigations centered on understanding how to best render these primary elements. Through a series of psychophysically inspired usability studies, we have determined that the functional line width for supporting vibrotactile information extraction should be 4mm and the minimal line width should be 1mm to support detection of vibrotactile elements. We note the distinction here between detection and extraction. If one simply wants to detect that a line (or point) is present, it need only be 1mm in size. If, however, a task needs to be completed that requires interpretation of the line, it needs to be 4mm in width. We have also established that the appropriate gap spacing between vibrotactile lines should be at least 4mm, which ensures that each element can be identified as a distinct entity [55].

While 4mm ensures discrimination of vibrotactile lines rendered parallel to each other, they are not generalizable to oriented lines that subtend an angle between each other. In Reference [55], we identified that a minimum cord length (and by extension the angular separation between two
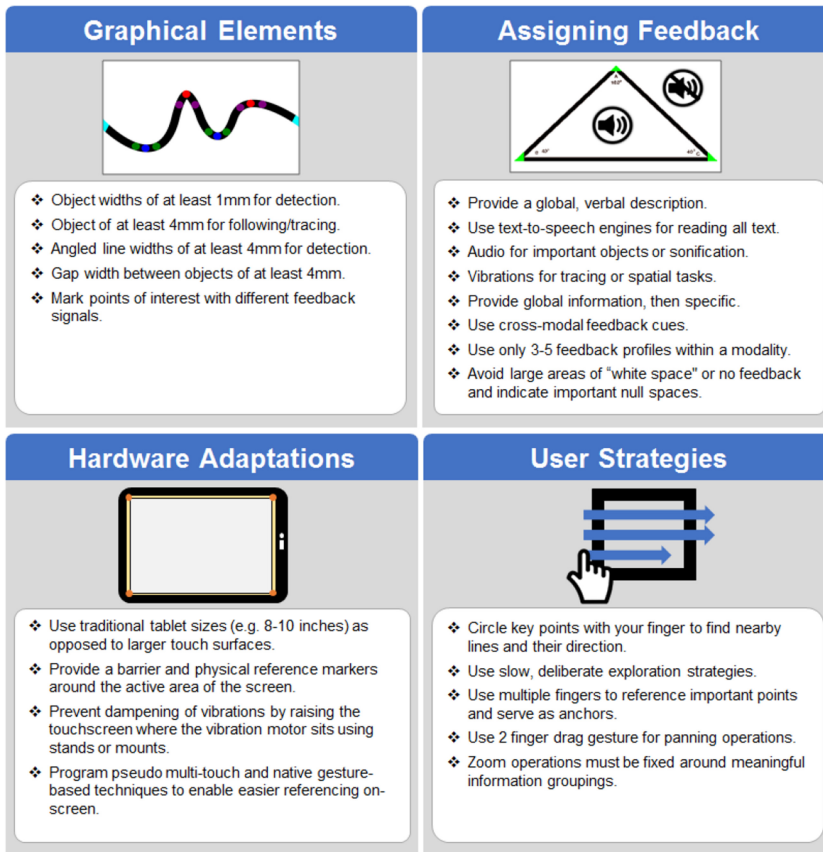
Graphical Elements

- ❖ Object widths of at least 1mm for detection.
- ❖ Object of at least 4mm for following/tracing.
- ❖ Angled line widths of at least 4mm for detection.
- ❖ Gap width between objects of at least 4mm.
- ❖ Mark points of interest with different feedback signals.

Assigning Feedback

- ❖ Provide a global, verbal description.
- ❖ Use text-to-speech engines for reading all text.
- ❖ Audio for important objects or sonification.
- ❖ Vibrations for tracing or spatial tasks.
- ❖ Provide global information, then specific.
- ❖ Use cross-modal feedback cues.
- ❖ Use only 3-5 feedback profiles within a modality.
- ❖ Avoid large areas of "white space" or no feedback and indicate important null spaces.

Hardware Adaptations

- ❖ Use traditional tablet sizes (e.g. 8-10 inches) as opposed to larger touch surfaces.
- ❖ Provide a barrier and physical reference markers around the active area of the screen.
- ❖ Prevent dampening of vibrations by raising the touchscreen where the vibration motor sits using stands or mounts.
- ❖ Program pseudo multi-touch and native gesture-based techniques to enable easier referencing on-screen.

User Strategies

- ❖ Circle key points with your finger to find nearby lines and their direction.
- ❖ Use slow, deliberate exploration strategies.
- ❖ Use multiple fingers to reference important points and serve as anchors.
- ❖ Use 2 finger drag gesture for panning operations.
- ❖ Zoom operations must be fixed around meaningful information groupings.

Fig. 1. Our recommendations and considerations for creating multimodal graphics on touchscreen platforms.

oriented lines) of $4mm$ should be maintained for accurate detection and discrimination of oriented vibrotactile lines.

We have also conducted explorations into line following and line tracing. Such tasks are important toward a better understanding of how individuals might be able to recognize outlines of important features or to follow important graphical elements (such as trend lines or bars). In Reference [73], we investigated path following and recreation of both linear and nonlinear profiles with 21 individuals with VI. We found that providing borders of a differing feedback (auditory borders with a vibratory line or vice versa) are helpful in supporting fine-tracing tasks, but single modality lines are sufficient for most other tasks. If a profile is nonlinear, we found that marking notable points (e.g., inflection points or start/end points) with additional feedback significantly increases the accuracy of being able to follow such profiles. In the study we conducted, we used vibration to represent the line and auditory tones to represent inflection points. In this work, we also investigated whether a single inflection point or multiple points varying in tone were more successful in illustrating impending changes in direction. From a performance standpoint, both the single and multi-point approach were successful, with all line profiles being traced with average deviations between $7-9mm$ and no statistically significant differences found between the two cases. However, 33% of users preferred the multiple tone case and 22% of users preferred the single tone

Table 1. Graphical Elements Recommendations

Graphical Elements
1. Object (e.g. line) widths of at least 1mm for detection.
2. Object (e.g. line) widths of at least 4mm for following/tracing.
3. Angled line widths of at least 4mm for detection.
4. Gap width between objects of at least 4mm.
5. Mark points of interest (e.g. end-points, inflection points, vertices) with different feedback signals.

case—illustrating the importance of flexible design and selectable UI parameters in such displays to accommodate highly varying user preferences. The concept of using feedback across modalities to help differentiate important points is supported by other research, including References [24, 60].

Extending this notion even further to basic geometrical shapes, we investigated the identification of basic shapes (up to five vertices) on a vibratory touchscreen and found that switching vibratory signals (e.g., increasing the intensity of the vibration) at the vertex points aids in shape identification accuracy [72]. This change in vibration patterns gives users key points to identify and salient landmarks to count and to return to during their exploration. We note that conveying more than four vertex points in this manner makes the task overly complex. A large number of vertices seems to cause users to lose their place or to lose track of previously accessed points, thus limiting the ability to accurately identify the shape. This work illustrates that signaling can also be done within modality (as was done in this case by using different vibration patterns), but this comes with limitations. We encourage designers to use multimodal cues when possible to alleviate the cognitive workload of having to differentiate within limited bandwidth modalities (e.g., vibrotactile perception).

A summary of the optimal rendering guidelines for points and lines is shown in Figure 1. The premise of these fundamental perceptual studies is to push the community toward adopting empirically validated and perceptually rooted standards that inform the creation of these non-visual, vibrotactile graphics, and we encourage further work in establishing such guidelines that bridge multiple domains toward informing multimodal graphical representation at large.

Collectively, our group recommends the following perceptually rooted design guidelines for creating multimodal graphics in Table 1.

3.2 Feedback Assignment

In addition to sizing elements, questions also arise about what feedback profiles to use and where to use them. Here, we detail which vibration and auditory profiles have stood out as being the most easily identifiable and distinguishable in the context of continuous exploration of graphical entities. From an auditory perspective, simple tones tend to be preferred over highly stimulating profiles. There is a wealth of information available on using sonification techniques successfully in relaying graphical information, and we refer users to References [23, 48, 49, 59, 79–81] for more details. In the context of representing graphical elements, we tend to use tones to mark important points or to signify gradients or changes in data.

We also note that all text displayed within a graphic should be read aloud via text-to-speech (TTS) engines. This work leverages a built-in TTS engine to read text aloud, as a pillar of this

work is focused on using readily available hardware/software to minimize the barriers to adoption. It is acknowledged that human recorded speech may be preferred over TTS engines and that including multiple concurrent voices can improve the speed and comprehension of scanning for information [28]. These alternatives could be considered in future work to provide options for personal preference. Incorporating TTS can be done in multiple ways. Currently, we have found it best to have text read aloud when a user's finger is over the text in the diagram. We also recommend a gesture-driven text description of the context of the image that can be accessed at any time. This description provides an overview of the graphic (similar to a caption or alt tag) to the user for providing context and allowing for invoking a mental schema of that type of graphic.

For multimodal graphics, we tend to use vibrations to represent lines or objects that require tracing. There has not been as much investigation on which vibration profiles are most conducive for continuous exploration in the context of graphics. There is, however, rich literature on discrimination thresholds and subjective interpretations of vibration profiles more broadly [14]. If the frequency of a vibration is below 3 Hz, it is perceived as slow kinesthetic motion [68]. From 10 to 70 Hz, one feels rough motion or fluttering, and from 100 to 300 Hz, smooth vibration [68]. In the perceptual space of vibrotactile stimuli, such qualitative differences are expressed by two distinct perceptual axes composed of 40–100 Hz vibrations (categorized as smooth and clear) and 100–250 Hz vibrations (categorized as jagged and bumpy), respectively [33]. Our group has begun to investigate which vibration profiles are most optimal in the context of the operating range of the actuators embedded within touchscreens and which signals best pair with traditional tactile line profiles. Generally, we recommend using vibration feedback with profiles that offer subtle, low-amplitude sinusoidal behavior, as opposed to high-intensity, constant buzzing for extended lengths of time. The latter can often increase the risk for sensory fatigue in the fingertips [60]. In a recent work [74], we investigated which vibratory effects best mapped to embossed line profiles. In this work, we found vibrations at 2.5 Hz best mapped to embossed dashed lines, 10 Hz to embossed dotted lines, and 50 Hz to embossed solid lines. We also note that it is not recommended to use more than three to five vibration profiles total. In our collective studies to date, we have not found a need for more than this, and more importantly, acknowledge the ineffectiveness of using too many signals of the same modality from a perceptual and cognitive standpoint.

Instead, we promote the use of multimodal feedback strategically. We have found that vibration cuing tends to be more effective for spatial information, whereas speech and auditory cuing for semantic information. More importantly, the notion of information layering is important given low resolution and limited real-estate of the hardware platform. We encourage designers to provide a summary verbal overview of the image to provide the global context of what is being displayed on screen, and then methods by which users can "drill-down" into the image and its components. This drill-down is often done via gestures where users can move "in" and "out" of different layers of information. Excellent examples of this technique have been demonstrated in References [61–63].

Finally, there is a balance of feedback and no feedback that must be provided. Too much information conveyed on the screen results in confusion, and too little makes navigation and interpretation challenging. Further, what we deem as the "Lost in Touchscreen Space" problem is complementary to this. Essentially, this problem relates to areas of the screen that provide no feedback or is sometimes coined "negative space." We have anecdotally observed that users who are exploring onscreen content often look for some type of feedback semi-frequently, otherwise concerns arise of having drifted off of the object of interest or out of the active screen areas. To address the "negative or white space" problem, we have found that providing cues (a low-level clicking sound or even a verbal cue stating "no information") is preferred rather than leaving a lot of open, white space that does not provide any feedback.

Table 2. Feedback Assignment Recommendations

Feedback Assignment
1. Provide a global, verbal description of the graphic.
2. Use text-to-speech engines for reading all text aloud.
3. Use auditory feedback (e.g. musical tones) for marking important objects or for implementing sonification techniques.
4. Use vibratory feedback for objects requiring tracing, following, or spatial relations.
5. Layer information such that global information is provided first, followed by opportunities to "drill down" into additional information.
6. Use cross-modal feedback cues whenever possible to alleviate cognitive workload.
7. Avoid using more than 3–5 feedback profiles within a modality.
8. Avoid large amounts of "white space" or areas of no feedback and consider indicating important null spaces (such as the inside of a shape).

Collectively, our group recommends the following suggestions for assigning auditory and vibratory feedback in rendering multimodal graphics in Table 2.

3.3 Hardware Adaptations

Some of the most obvious challenges to touchscreen-based graphics are hardware limitations. First, vibration feedback is not localized, resulting in the majority of the exploration needing to be done via a single finger. This limits one's ability to quickly capture a global view of the on-screen content. As a result, traditional exploration strategies using multiple fingers or hands as a frame of reference are not yet technically possible, though interesting work affixing external vibrators to fingers has been investigated (e.g., Reference [2]). Only the center of the finger's contact pad with the touchscreen is used to generate the on/off vibratory feedback. This can be a problem, depending on how the user's movement intersects the stimulus, and is particularly challenging with small stimuli. This single touch limitation makes detection of angled or curving trajectories particularly challenging, relying largely on the kinesthetic system, which is rather noisy. As a consequence, while people may be able to trace these stimuli, it is often difficult for them to accurately perceive the magnitude of the deviation.

In our work, we have found two primary work-arounds to address the single-touch hardware limitation. The first is that a pseudo-multitouch can be programmed into some commercial Android touchscreen phones and tablets to account for multitouch in a narrow way. This can be done by detecting and registering the number of fingers that are present on the screen as recorded chronologically. The first finger to make contact with the screen can be designated as the "active finger," triggering the feedback effects on the screen. While other fingers cannot trigger feedback, they can act as guides and reference points to remember key places on the screen. These reference fingers, however, should not come too close together, as a new centroid point would then be calculated between the two fingers and could result in spurious triggering of feedback. From a practical implementation, we have found this process to be most useful in teaching individuals to use complementary fingers as anchoring and referencing points. For example, users may anchor other fingers of their tracing hand (such as the thumb) to the sides of the screen to assist their tracing of

a line on a screen. The aforementioned evaluations on non-visual panning and zooming techniques (i.e., two-finger panning and functional zooming) were also based on such pseudo-multitouch techniques. We note that these reference fingers need not all be used on a single hand, and in fact, using two hands, is sometimes easier, particularly if the user is referencing edges of the screen. The above techniques are best suited for traditional tablet sizes (≈ 9–$10''$). While larger touch surfaces have been investigated and multi-touch techniques proposed (e.g., Reference [38]), most of the larger tablets often do not contain vibrotactile feedback and are more expensive and less portable than traditional tablets. We note, however, that with the advent of Bluetooth vibration motors, as well as wearable mechanisms, the lack of vibration capabilities in larger screens may not necessarily be a problem if individuals are open to augmenting their experience with additional hardware.

Given the limited resolution of vibrotactile perception and the small size of most touchscreens, it is important to maximize the information content that can be displayed via the limited screen real estate. This does not necessarily indicate that large touchscreens have inherent benefit over small touchscreens, such as mobile phones. In fact, a recent investigation from our group comparing screen size effects on a pattern matching task demonstrated that compacting graphical features and making them smaller to fit a mobile phone display does not necessarily negatively impact a user's interpretation of the graphic [71]. On a mobile phone, a simple graphic can often be identified much quicker than on a tablet, given the smaller amount of space that must be explored. There is certainly a limit to the amount of information that can be conveyed on smaller screens, though layering and panning techniques will likely help with this. We encourage designers to keep in mind the complexity of the graphic being displayed and the intended outcomes of the exploration (accuracy vs. time) when choosing the platform to display such graphics. Nonetheless, we encourage sizing the graphic to fit the screen. Often, for traditional mobile platforms and tablets less than 10 inches, this means filling the screen area, but for larger tablets, this notion is not necessarily upheld.

Further, the lack of physical boundaries on the active screen area of a touchscreen make it very difficult to know whether one is exploring in the active working area of the screen or if one is on the borders (see Reference [15] for a brief discussion). To address this, we have used easy physical adaptations to the touchscreen itself (see Figure 2). We have found that adding small, rubber dots along the vertical sides of the screen, or placing thin plastic tape across the perimeter of the active screen area, helps users in both ensuring that they remain in the active screen area and that they can also access easily identifiable reference points. This echoes recommendations put forward by Reference [15] to assist users in staying within the active area boundaries of a touchscreen through the use of simple hardware modifications, such as stickers. To prevent vibration dampening from vibrations dissipating from the device to other surfaces, we recommend an additional hardware modification—affixing stands or mounts to the bottom of the tablet when possible (see Figure 2, back and side view). While we have found that the reference markers and pop-sockets support comfortable ease of use in exploring the tablet on a table-top or desk environment, alternative poses such as holding the tablet in one hand and exploring with the other are also viable options.

We have also employed the use of gestures when possible (in line with built-in accessibility gestures of the native operating system) to aid in overall operation and navigation on-screen. We believe that a dual-pronged approach of establishing a framework for creating content and developing guidelines for exploring them will lead to the best chances of success in the digital, multimodal graphics space.

Finally, all of the above studies have been conducted on Android-powered hardware platforms. This was done for several reasons. First, companies that use the Android OS, such as Samsung and LG, have and continue to consistently have vibration capabilities across mobile and tablet platforms. While Apple has had touch capabilities in their mobile platforms for some time, it has yet to be available on tablets. Second, Android provides significantly more flexibility in developing

Fig. 2. Front, back, and side view illustration of potential tablet adaptations. Plastic tape sets the boundary of the active area of the touchscreen. Rubber stickers provide quick reference as to the location of the four corners of the active area. Rubber stickers are also used on back to prevent the tablet from moving on a surface during exploration. Collapsible stands lift the tablet off of the surface at an angle to prevent dampening of tablet vibrations and to provide a comfortable angle for exploration.

vibration profiles and providing continuous vibration feedback. To date, Apple has given limited control over the tactile feedback capabilities and limits the run-time for them. Apple has recently released Core Haptics for newer iOS devices (iPhone 8 and above), which promises to provide more flexibility to developers using tactile feedback in their applications [16]. Given that majority of VI users have adopted iOS, this movement is promising and necessary in supporting multimodal user experiences, learning, and inclusive design more broadly.

Collectively, relatively simple adaptations can be made in hardware and software that make exploration of multimodal graphics easier. A summary of our recommendations can be found in Table 3.

3.4 User Strategies

While authoring multimodal graphics, we have discovered the importance of being cognizant of the user's approach to interpreting them. When individuals learn tactile graphics, they are taught certain strategies to help maximize information encoding and interpretation of the information through touch. Likewise, we have found new users of multimodal graphics may need some training on how to best perceive and extract vibrotactile information. Although every interface will be different, the hand movements (and thus encoding strategies) used to accurately perceive and extract vibrotactile data will likely look similar to the ones we describe below. While our group has not conducted formal studies comparing one exploration technique from the other, we have informally aggregated common exploration strategies that have resulted in high achieving outcomes from previous user studies. We have included these observations within our guidelines

Table 3. Hardware Adaptation Recommendations

Hardware Adaptations
1. Use traditional tablet sizes (e.g. 8–10 inches) as opposed to larger touch surfaces.
2. Provide a barrier and physical reference markers around the active area of the screen.
3. Prop up the end of the touchscreen where the vibration motor sits (e.g. stands or mounts) in order to get the full effect of the vibration feedback.
4. Consider programming psuedo multi-touch techniques and native gesture-based techniques to enable easier referencing and navigation on screen.

to encourage designers and developers to incorporate some common strategies that we have identified as being common to and beneficial for VI users.

With physical media, there are three main exploratory procedures users employ to access and extract graphical information: (1) lateral scanning, (2) edge following, and (3) global exploration with the whole hand [36, 41, 44, 50]. Even with physical media, this way of accessing visual information is very cognitively demanding due to the integration of spatial information across time and space. This integration is also not always precise, which can negatively affect creation of a mental image with accurate spatial features [83].

By contrast, touching a graphical line on a touchscreen device simply leads to the sensation of vibration suggesting that you are on the line. Determining the line's orientation requires active finger movement, but lacks directional information. Some successful strategies we have observed users employing on vibrotactile displays to accurately determine a line's orientation include (see Figure 3 for illustration): circling a part of the line, usually at a junction, to determine its orientation and the number and direction of intersecting lines (A), zig-zagging along the line to follow its profile (D), and using multiple fingers as reference points on the edge of the screen to remember important parts of the line, such as the start and end points (E) [60, 72, 73]. Unlike traditional hardcopy graphics, touchscreens have the advantage of using multiple modalities. Marking inflection points with auditory feedback has also been successful in letting users know that there is a change in orientation of the profile [73]. This approach carries over into shape exploration, where marking vertices with a different modality of feedback aids in shape identification [24, 72]. We have also aggregated shape and line exploration strategies observed across multiple experiments and found that users employed additional techniques including making crosses at vertex points to find connecting line segments (B) and scanning horizontally across the shape to uncover symmetry (C) [29, 72].

Given the limited screen real estate on commercial touchscreen devices, accessing graphical materials such as maps would not be possible without the incorporation of panning and zooming operations. Towards this end, Giudice and Palani have done extensive work on navigation in the context of maps, particularly in the areas of how to incorporate zooming and panning operations for non-visual vibrotactile access. By comparing map learning performance with four different panning techniques across cognitive and usability measures, we have found that a two-finger drag technique is best served for performing nonvisual panning operations. With this technique, the user can initiate "pan mode" by placing a second finger on the screen and dragging the underlying map synchronously. Upon removal of the second finger, the interface resumes to exploration mode from the same location on the map, thus allowing the user to access the graphical elements after panning from a known reference [53]. In a follow-up study, we evaluated users' ability to perform nonvisual zooming operations that support information integration across zoom levels.
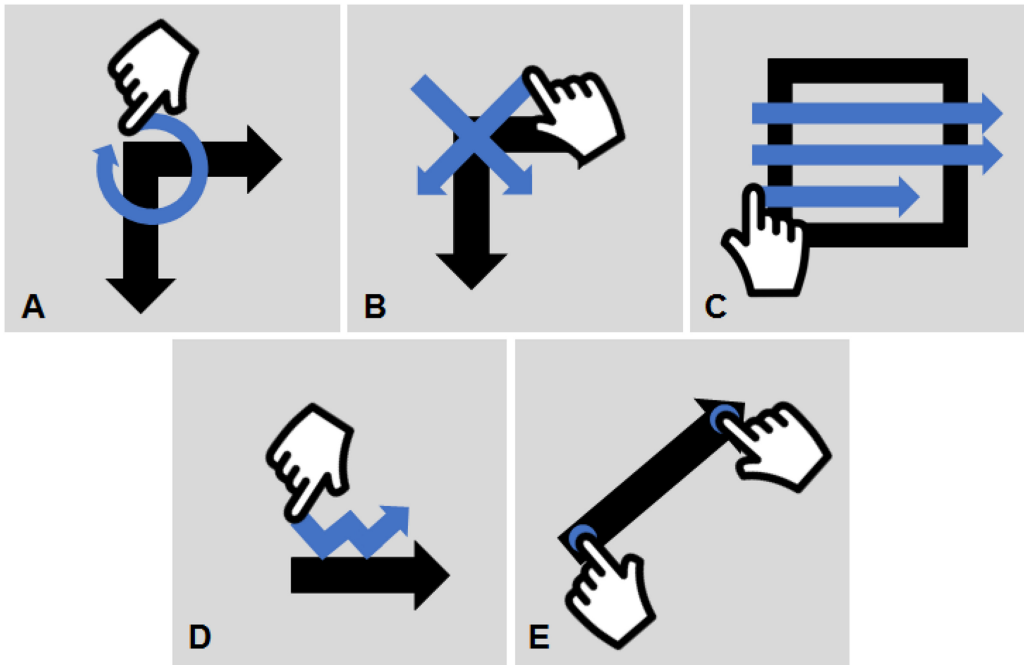
Fig. 3. An illustration of the five most common strategies for nonvisual exploration of multimodal images on a touchscreen. (A) Circling a vertex point; (B) Crossing a vertex point; (C) Horizontal (and vertical) panning; (D) Zig-zagging to follow a line profile; and (E) Using multiple fingers to reference points.

Table 4. User Strategy Recommendations

User Strategies
1. Circle key points with your finger to find nearby lines and their direction.
2. Use slow, deliberate exploration strategies.
3. Use multiple fingers to reference important points and serve as anchors.
4. Use 2 finger drag gesture for panning operations.
5. Zoom operations must be fixed around meaningful information groupings.

We compared two different zooming methods (i.e., a fixed-zoom, where levels were based on the scale of the map and a functional-zoom, where levels were based on meaningful information grouping). Results found that across cognitive and usability measures, the functional-zoom method was most efficient for supporting non-visual zooming operations, which was congruent with earlier work that evaluated level-based zooming operations on tactile diagrams [61, 63]. Findings also suggested that users were able to accurately integrate and relate graphical elements even when presented across zoom levels [54].

Collectively, our recommendations for successful exploration strategies of touchscreen-based multimodal graphics can be found in Table 4.

The guidelines listed in Table 4 are design recommendations meant to serve two purposes: (1) promote consistent display of multimodal graphics and (2) motivate researchers to consider

innovative ways to create multimodal graphics that embed these design recommendations. Efforts to create multimodal graphics are already underway, with initiatives like Vital's Content Creator demonstrating the creation pipeline of such images [78]. We note that finding the appropriate balance of manual control and automated assignment is an interesting research challenge that we are continuing to investigate. These guidelines serve as a basis from which such creation frameworks can be established.

4 ITERATIVE DESIGN AND PROTOTYPING OF MULTIMODAL GRAPHICS

Using the guidelines and strategies above, our team has begun designing multimodal graphics, particularly in the area of mathematics and statistics, which are known to have highly visual content that is often difficult for individuals with VI to fully participate. We have focused on two areas that extend from our fundamental explorations: geometry and bar charts. From a teaching and learning perspective, bar charts are regularly used to represent categorical datasets and are considered a benchmark in fundamental statistical analysis. These types of graphs provide many of the fundamental aspects of all statistical graphics, including labels, legend, axes, and increments. From a graphical perspective, students must be able to use both axes to determine relevant information while surveying the individual data provided as "bars." These skills are fundamental to utilize more advanced graphical representations of data (line plots, scatter plots, etc.). Geometry figures are also fundamental to larger STEM concepts. From early grades, students are taught the parts of basic shapes beginning with points (i.e., nodes) and sides. As they progress forward, concepts of angles, relationships between angles, and angles with sides are explored. This knowledge is the backbone for advanced geometry and trigonometry, which has many connections to physics. Therefore, these two concepts provide fruitful ground for scaling up the design guidelines and strategies employed to date but also represent fundamental concepts in mathematics from which additional complex graphics can be designed. The bar graph and geometric shapes were also selected, as they provided a balanced selection of graphics to be explored by individuals with VI for the purposes of this work. While the bar graph primarily relies on horizontal and vertical movements of the hands when exploring, the geometric shapes (triangles and trapezoids) were selected to add a level of complexity. The figures not only include diagonal lines but also include multiple intersecting lines, creation of multiple angles, and more complex relationships between parts of the graphic. Finally, these two concept areas are widely used in the classrooms, in the workplace, and in both scientific and popular press writing, enabling the findings of this work to have practical implications.

A representative example of two of the graphical concepts rendered are shown in Figure 4. This figure presents how a bar chart and geometric figure may be rendered multimodally. Feedback profiles of these renderings include easily distinguished vibration patterns (such as 10 Hz and 50 Hz; see Reference [74]) for each bar, text-to-speech for any labels, and for the geometry figure, sounds and tones where appropriate to highlight important information. Equally as important is the treatment of white space. Meaningful white space (such as to represent the inside of a triangle) must be marked and we have done so with soft, auditory clicking. Further discussion on feedback assignment for this figure can be found in Section 4.1.2.

4.1 Design Process

Our design process involved three distinct and important steps: (1) Establish the pedagogical goals of the graphic; (2) Determine the pedagogical value of feedback assignment; and (3) Prototype, validate, and iterate the graphic. These steps streamline the process of creating multimodal graphics with appropriate and accessible audio and haptic functionality.
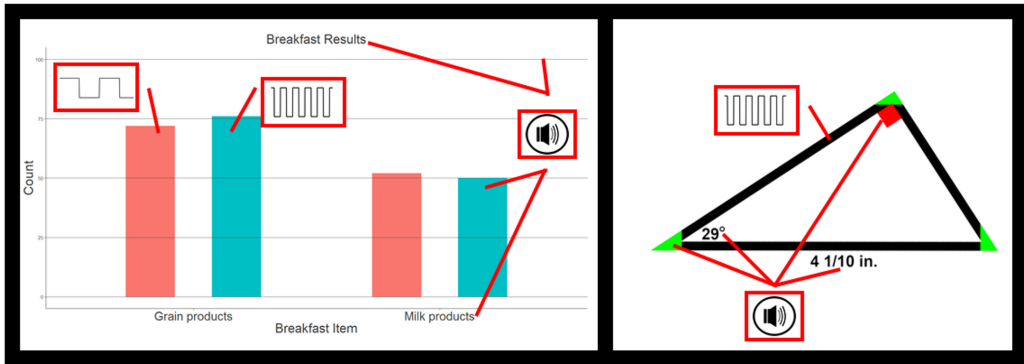
Fig. 4. An example of a bar chart (left) and geometric figure (right) rendered digitally with multimodal feedback.

The following sections illustrate each of the three steps outlined above for the creation of both bar charts and geometry figures, examples of which are given in Figure 4.

4.1.1 Establish Pedagogical Goals. The first step in designing accessible, multimodal, non-visual representations of graphics begins with establishing the pedagogical goals of the graphics. The definition of these goals will dictate where audio and haptic effects are placed and what function they serve within the graphic. It also dictates how much information is outright given versus what must be extracted from the user.

We chose bar charts that had no more than four categories, and we chose various triangles and quadrilaterals with angle, vertex, and length information available, as these represented typical images that would be used when introducing such concepts in an educational setting. Moreover, these relatively simple graphics provided an appropriate foundation to investigate multimodal graphic rendering more broadly.

Questions were created about each image that were (1) pedagogically meaningful and had embedded spatial tasks and (2) required individuals to extract information from the graphic and/or make comparisons about the data extracted. Examples of questions are found in Table 8.

4.1.2 Feedback Assignment. Perhaps one of the most challenging parts of creating multimodal graphics is determining what feedback signals should be assigned to what attributes in the image. It is important to note that this feedback must not only be appropriate for the individual elements of the graphic, but it must also contribute to the greater understanding of the image as a whole and appropriately serve the pedagogical purpose of the graphic. To better illustrate how the guidelines from Section 3 can be put into practice, we describe the feedback assigned to a bar chart and geometry figure of our own.

We begin with the bar chart shown in Figure 5. Before the user encounters any feedback, it is necessary to establish the context of this graph. This is done via a textual overview description of the image. Context is necessary for understanding the purpose of the content displayed onscreen and better prepares the user to perform any tasks with the data. This is a common practice with tactile graphics and is similar to the use of Alt Text or captioning in digital figures [8]. The text description of the graph in Figure 5 reads as follows, "We asked 2 questions to voters in Maine. One question was if they went out to vote in 2016 and the second question was if they planned to vote in 2020. They responded yes or no. For each question, we've recorded the percentage of yes and no responses for both questions in the form of bars."
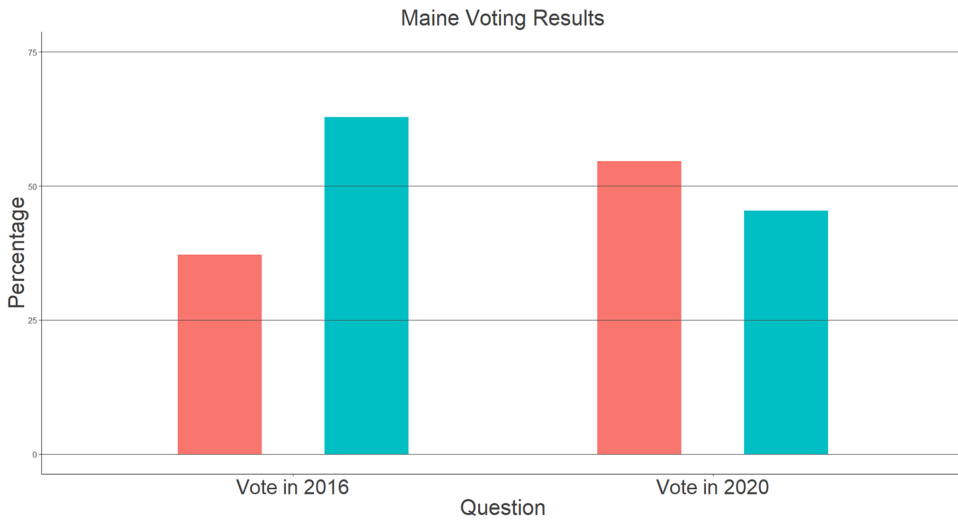
Fig. 5. A bar chart representing fictional voting behaviors of voters in Maine during the 2016 and 2020 elections.

We then begin by outlining all components that would be immediately served with audio or text descriptions. All text in the image is read aloud to the user when their finger is in contact with any part of the text string. For example, when a user touches the text area of the title, "Maine Voting Results" is read aloud. The same is true for all axes and data category labels. When a user's finger touches a bar area, the category of the bar is read aloud. For instance, when the user touches the right, blue bar in the "Vote in 2016" category, it reads "No, Vote in 2016." Having all labels read aloud is an important advantage provided by the touchscreen medium. Often, labels on embossed graphics have to be shortened or excluded from a bar chart due to space concerns [18]. To avoid audio fatigue from the user while tracing within a bar, the auditory label for the bar does not repeat until the user moves out of the bar into another bar, to accommodate back-and-forth exploration of the bar. If the user needs the label read again, they can also double tap on the bar to receive the audio label (as well as the y-axis value).

In feedback assignment for the y-axis gridlines, our team decided to have the values of the gridlines read aloud when the user's finger crossed them. For instance, when the user crosses the "25" gridline, the value "25%" is announced. This choice was made to eliminate the need for users to trace from the bars back to the axes labels to garner or compare information between the two zones. This decision was in line with the goal of this graphic (information extraction) in our particular case. However, should the goal of the graph be for users to use the y-axis more extensively, users should be required to trace back to the axis, and thus gridlines should not announce their values to the user.

After establishing all components served by audio, we branch into haptics. In this bar chart, vibration is determined to be the most helpful for specifying the bars. The bars themselves can be thought of to represent physical or textured objects. We know from Section 3 that vibrations can be used to make these types of components more salient through touch. Not only do the bars announce their labels, they also apply a constant vibration to the user's fingertip when touching the bar area. To differentiate the data categories, red bars have a different vibration profile compared to blue bars. We chose patterns that were more likely to be differentiated, such as 25 Hz and 50 Hz (see Reference [74]), to improve the user's ability to distinguish the two different data categories.

Fig. 6. Trapezoid ZWXY with the lengths of three sides labeled and no angles labeled. Black line segments vibrate; red squares indicate right angles and read aloud "90 degrees" when a user encounters them; and green areas behave similarly to right angles except that they are not labeled and therefore only play a tone to indicate that they are angles.

Vibration occurs only within the bar, to allow the user to explore and trace within that context, as is helpful. When a user's finger is on the bar, they feel a specified vibration, and when they are off of the bar, they feel no vibration.

Finally, we end the design considerations with enhancements from the technology itself. It is important to note that there is no legend displayed on the bar chart. This is a distinct advantage of the touchscreen technology. Information can be layered, rendering some visual tools such as legends, which take up space and require moving haptic exploration away from the core stimuli, unnecessary. Instead, this omitted, visual information can be embedded within the graphic itself for ease of exploration, as is the case in our bar chart. Users are also afforded typical touchscreen gestures, such as swiping and tapping, which may assist them in comprehending the graphic. Implemented in this graph is a double-tap feature that users can utilize anywhere on the graph to be verbally given the y-value of the chart at that particular point.

Although a challenge in some graphical representations, the white space in this graph provided no feedback to the user. Not all white space must be indicated with feedback. If the white space is not meaningful to the graphic, it is extraneous information that may over-saturate the user during exploration. If the white space is not meaningful to the context of the graphic, and if there is already sufficient information to derive a conceptual overview of the graphical entity from other context and feedback modalities, white space does not need to indicated to the user.

Next, we use one of our geometrical figures (Figure 6) to illustrate how the guidelines from Section 3 can be used to create a multimodal representation of a standard mathematical figure. Similar to the process outlined above for the bar graphs, context is first determined to provide the user with information about the image. This verbal description states that "This is a geometry figure. Geometry is a way to refer to shapes, like squares, circles, and triangles."

We then assign audio feedback. Like the bar graph, text should be read aloud wherever text labels appeared on the figure. This extends to any explicitly mentioned angle measurements, segment lengths, and named segments or angles. To assist the user in comprehending the graphic, we have found during discussions and focus groups with individuals with VI that it is better to include more information in the "summary description" beyond simple names and measurements. Segment

directions (horizontal, diagonal, and vertical) and more comprehensive angle names (angle ZYX instead of angle Y in Figure 6) can be added to the audio label as well. To distinguish important components of the shape, such as corners or right angles, a short audio effect can be used as an indicator for these elements. In our implementation, regular angles were associated with a brief "ding" effect whereas right angles announced their measurement, "90 Degrees."

Vibrotactile feedback is used to indicate line segments. Just as vibrations can assist in the physical interpretation of bars, they can be used to convey physicality of lines in geometric figures. In a recent study [74], we found users were able to best distinguish the following vibration profiles in terms of traditionally embossed lines:

(1) 2.5 Hz to represent slow or dashed sensations;
(2) 10 Hz to represent quick or dotted sensations;
(3) 50 Hz to represent constant or solid sensations.

Each vibration pattern in Android consists of alternating off-on calls in milliseconds to the device's motor. For instance, we used a constant, low-velocity buzzing sensation (pattern: [100, **50**, 100, **50**, 100, **50**], where a **bold** number corresponds to the duration the motor is on) to create a low-level, salient, representation of a line in Figure 6. In this pattern, the first number, 100, corresponds to 100 ms of motor off time. The 50 ms following it is the duration in which the motor is on. This off-on pattern repeats indefinitely, creating a low-velocity buzz. In addition, other vibration patterns were used to add emphasis to important parts of graphics, such as for shape vertex points (e.g., pattern: [25, **75**, 25, **75**, 400, **50**, 200, **0**]).

A double-tap feature was added to the geometry graphics to provide the user with information about the line or angle over which the double-tap occurred. For instance, a user may wish to know what line they are currently tracing and can double-tap on the line to re-verbalize the line information. We were careful in our implementation to avoid complicated gestures, such as those that are multi-stroke or multi-finger gestures. We have observed (along with References [13, 15, 27]) that executing these types of gestures can be difficult for a VI user to perform due to the timing and spatial constraints of the gesture. In the case of multimodal graphics, the pool of available gestures is further limited. One-finger gestures, such as swipe, drag, and single-tap, are often disabled to give users the freedom to explore graphics by freely moving one finger around the screen. Therefore, even though they may be more difficult to use, it may sometimes be necessary to include double-tap gestures. However, due to the difficulty in timing and positioning multiple taps with multiple fingers for VI users, we do not recommend going beyond a double-tap gesture and to keep the number of fingers required to double-tap as low as possible.

Unlike in the bar chart where there was significantly more information and context to act as sufficient points of reference, the geometry figure has much more white space. For bar charts, white space and the absence of feedback always represented the same thing—no data present in this area. For hollowed out, outlined shapes, such as our figure, white space has two distinct roles: One, to let the user know there is nothing of interest in the area; and two, to let the user know they are inside the bounded polygon, the latter of which affects the understanding and exploration of the shape itself. To mitigate the challenges associated with white space, a clicking sound can be added to meaningful white space, inside of the shape. This indicates to the user that, while there is no data where they are exploring, this white space is still relevant to the perception of the image as a whole. This recommendation was inspired by a feature found in VoiceOver-enabled iOS devices—when in a menu and there is nothing to select, a constant clicking noise can be heard.

The bar chart and geometry figures illustrate practical application of the design guidelines while also highlighting the challenge of broad application of guidelines across images. Many of the strategies and design decisions made were done so after significant iteration and trial-and-error and

were largely driven by context and task goals. It is important to interpret the guidelines posited in this article in this context, acknowledging that graphics and their interpretation are largely open-ended. While guidelines that support detection and extraction from a perceptual standpoint are necessary, there is an additional level of complexity driven by pedagogical goals that must be considered in rendering multimodal graphics. Regardless, both examples illustrate how text, sound, and vibrations can be used cohesively to represent basic graphics nonvisually via the touchscreen of smart phones and tablets.

4.1.3 Software Prototyping. Multimodal images are displayed via an Android application. Using Google Play Store's capability to send development versions of our Android application, we iterated on graphic design and feedback assignment on a weekly basis for six months. We started by integrating haptic feedback and text-to-speech on a single image using a single screen application. To allow for easier control over the vibration system in the application, we integrated our previously developed, open source Android haptic library, which comes with established vibration patterns and higher-level control over the Android vibration API. For text-to-speech capabilities, we used Android's native text-to-speech API.

The first iterations of the application used differently colored areas to trigger specific feedback. Evaluating the color beneath the user's finger is one of the easiest and quickest ways to implement multimodal prototypes in Android. This approach, however, has some limitations. First, every image must have its distinct components differentiated using a graphics program. Second, depending on the device, specific color values are evaluated differently. For instance, although a shade of red has not been altered in a graphics program, Android may read this value differently across devices affecting the operation of feedback. Third, it does not support more complicated interactions, such as overlapping areas.

Thus, we replaced the color method with an overlay of invisible polygons, which constitute triggers for the different interaction events. These polygons were defined by lists of points. However, in the future, we will implement a more streamlined approach—combining the images with their polygon definitions into one file format for organization and ease of use in making retroactive updates to the image.

To trigger feedback, a user's touch will cause the application to check if the touch was in a predefined, polygonal area based on a priority system. Areas with the highest priority are checked first. Then the system causes an event to occur based on which area was touched. The event will trigger text-to-speech, sound cues, or vibrations, depending on what is predefined for that area. An example for an area defined in this way might be a bar in a bar graph. The rectangular polygon around the bar defines the touch area in which the first touch will trigger the name of the area to be read and start a continuous vibration while the user keeps touching the same area. When leaving the area, the vibration stops, but the system remembers the text of the last area being touched and does not read its name again until another area has been touched. This is to prevent sensory fatigue in the audio space, as described in the section above.

A more complex interaction between areas comes into play when a horizontal grid line from the y-axis intersects with the bar. In this case the system will recognize that the user is touching the bar and execute the behavior tied to the bar, but will also check whether the user touched the grid line and perform the behavior for the grid line.

Grid lines were initially introduced in the project to give accurate information regarding the value of the y-axis to the user while exploring the bars. The grid lines read the corresponding value of the y-axis when they are touched. A double-tap gesture was also implemented to give the value of the y-axis at the point of the first tap. The double-tap gesture determines its position on

Table 5. Additional Themes That Arose during Software Prototyping and
Rendering of Multimodal Images

Additional Application Considerations

1. Consider providing labels and other key items during the contextual overview.
2. Avoid audio fatigue by reducing the number of repeat triggers.
3. Allow labels or important information to be read again via a simple gesture.
4. Assign feedback in a clear, consistent ways throughout the graphic.
5. Center and maximize the graphic on screen to support easy navigation and avoid the unintentional press of buttons or UI elements.

the screen relative to the position of the top grid line of the y-axis and the bottom one (where the bottom line represents zero on the y-axis) and then speaks the position's y-value, e.g., "65%."

After fundamental interactions were functioning, we scaled the application up to enable the capability to explore more than one graphic with the same interaction paradigm. Forward and backward buttons were implemented using the menu bar on the top of the screen to switch between images. It quickly became apparent, however, that this implementation reduced screen real estate and introduced accidental presses from users. As a result, we switched the menu bar out for a pull-out menu that is accessed via a three-finger swipe gesture. The slide out menu was later expanded to allow for custom orders of slides and recording of data for experimental testing. The purpose of this swipe menu was to assist experimenters in quickly switching through graphics and options related to each graphic and is not an implementation of a menu that we recommend when considering use by VI users (see Section 4.1.2 for further discussion on gestures).

Gestures are important interaction techniques that do not come at the cost of screen real estate or require audio or text input from the user. Whenever possible, user interface interactions should utilize gestures. We suggest using gestures to provide a simple way for users to access feedback options, to close windows, and to move to new activities. However, using gestures to achieve graphically motivated goals requires further investigation.

After applying our design guidelines across both bar graphs and geometry figures and prototyping and iterating our software application, a few additional themes important to consider at the application stage came to light. We suggest considering these additional recommendations, shown in Table 5, during the application development stage.

5 EXPERIMENT

To investigate the effectiveness of the guidelines proposed in the above sections, we conducted a within-subjects experiment comparing digitally rendered, multimodal graphics to their traditional, embossed analogs. In this experiment, participants were given three embossed graphics and their three multimodal counterparts on the touchscreen. Users were evaluated on each graphic with a series of questions.

The following study was approved by Saint Louis University's Institutional Review Board, and participants were provided informed consent and compensated for their time.

Table 6. Participant Summary

#	Age	Sex	Diagnosis	Age Diagnosed
1	25	F	Leber Congenital Amaurosis	< 1
2	63	M	Cataracts	33
3	22	F	Leber Congenital Amaurosis	< 1
4	67	F	Congenital Rubella	< 1
5	39	M	Retinitis Pigmentosa	5
6	48	F	Unknown	< 1
7	20	F	Glaucoma	< 1
8	36	M	Nerve Damage	2
9	31	M	Inherited Retinal Disease	< 1
10	37	F	Leber Congenital Amaurosis	< 1
11	51	F	Retinitis Pigmentosa	< 1
12	76	M	Cataracts	< 1
13	49	F	Retinitis Pigmentosa	7
14	45	M	Glaucoma	< 1
15	22	F	Ocular Albinism	< 1
16	42	F	Retinoblastoma	< 1
17	31	M	Nerve Damage	4
18	50	F	Optic Neuritis	8
19	47	M	Steven Johnson's Syndrome	7
20	37	F	Subdural Hematoma	2
21	21	M	Leber Congenital Amaurosis	< 1
22	52	F	Retinoblastoma	< 1

5.1 Demographics

Twenty-two participants with visual impairment participated in this experiment, recruited from the National Federation of the Blind (NFB) national convention in 2019. A majority of these participants ($N = 22$; average age = 41 years) were female ($N = 13$) and highly educated ($N = 16$ finished college or had graduate degrees). Participants represented a wide range of diagnoses, the most common being Leber Congenital Amerosis ($N = 4$), Glaucoma ($N = 3$), and Retinitis Pigmentosa ($N = 3$) (see Table 6 for complete demographics). Ten of the 22 participants stated that they had some form of light perception and all but 1 individual (who chose not to respond) stated that they had no residual vision and were completely blind.

To be considered for participation, individuals were required (by self-report) that they had sensation in their fingertips, functional movement of their fingers, no severe cognitive impairment (inability to understand directions or questions asked), and no/minimal hearing loss. Potential participants were also required to have the ability to read Braille and read tactile diagrams with minimal assistance. Of the 22 participants, 16 participants had sufficient experience with tactile diagrams to require almost no additional assistance, but 6 requested assistance reading the tactile diagrams.

To encourage participation at the convention, a $25 American Express gift card was offered to participants as compensation.

5.2 Procedure

Participants were evaluated on their performance interpreting bar charts and geometry figures (see Figure 7) on a tablet (9.7 Samsung Galaxy Tab S3) as well as on an embossed tactile graphic.
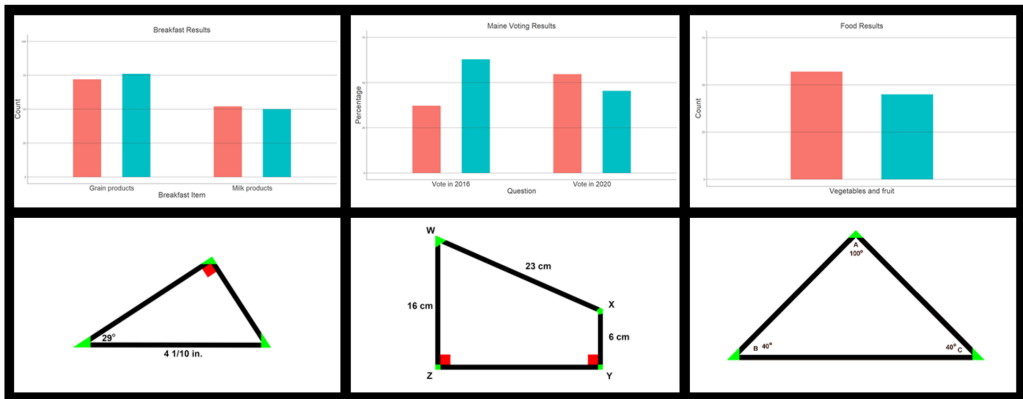
Fig. 7. Six images used in the experiment. The rightmost bar chart and geometric figure were used for training purposes. In the geometric figures, angles are denoted by a green color and right angles are denoted by red squares.

Each participant received the same three bar charts (one bar chart provided as training) and three geometric figures (one figure provided as training) rendered in both media. In practice, each participant was evaluated on six images total, with two being training images. A Latin square design was employed as a means to account for order effects during the study. Participants received all touchscreen or embossed versions of the images before the other version, with the order of bar charts and geometry images alternating according to typical Latin square methodology.

All hardcopy, embossed graphics were created with a ViewPlus Tiger EmPrint embosser in standard form representative of graphics given in an educational setting. The embossed graphics were comparable in size to the graphics given on the tablet to achieve mapping between the two stimuli and to eliminate any confusion moving between platforms.

Before receiving a bar chart or geometric figure (in either embossed or touchscreen form), participants were given time to train on a simpler graphic (see Figure 7, rightmost bar chart and geometric figure). Each training session lasted between 10 and 20 minutes and only concluded once participants successfully completed a set of criterion questions about the training graphic. During this time, participants were acquainted with the tablet's functionality, such as double-tapping to get more information on a bar or angle and learning how to differentiate between diagonal and horizontal lines. Some successful exploration strategies were provided to the participants, but they were also encouraged to discover other strategies that worked best for them. Participants were encouraged to ask questions during training or seek help finding an answer to a question, as they could not do so after this period in the official evaluation.

Once participants successfully answered all criterion questions, the official evaluation period began. This evaluation period consisted of two components: (1) exploration and (2) evaluation. During exploration, participants were given no time limit to establish the upper and lower limits for the task on both touchscreen and embossed forms. During evaluation, participants were asked four-to-six questions (see Table 8 for a set of questions regarding a bar chart and geometry image from the study) about the graphic that only involved information extraction; no questions asked for computation in any form. The standard number of questions asked about each graphic was six, with the exception of the questions pertaining to the triangle figure. For example, if a participant responded to the questions "What shape is this?" and "Are there any right angles?" with an incorrect answer, then they would not be asked the two follow-up questions associated with those questions. Depending on how a participant responded to a previous question, a follow-up question would not

Table 7. Performance Comparison

Embossed			
Category	N	M (%)	SD (%)
Overall	42	81.83	20.21
Geometry Only	22	85.53	14.80
Bar Chart Only	19	81.82	17.14
Touchscreen			
Category	N	M (%)	SD (%)
Overall	42	80.09	23.74
Geometry Only	22	86.44	18.44
Bar Chart Only	19	76.56	21.71

be asked and the minimum number of questions asked about the triangle would be four. Similar to exploration, evaluation had no time limit, and participants could take as long as necessary to answer the questions and could even return to the graphic for help answering the questions. Upon completion of the experiment, participants were thanked for their time and given their gift card.

6 RESULTS AND DISCUSSION

All questions used to evaluate participant knowledge of the graphics were scored and averages computed to summarize performance in the evaluation task. Twenty-two participants completed the geometry figure evaluations in both mediums, but only 19 participants completed the bar-chart evaluations, as 1 participant did not reach the bar charts in the time that was allotted. On average, a participant's performance on multimodal, touchscreen-based graphics ($M = 80.09\%$) was on par with their performance on the traditional embossed graphic ($M = 81.83\%$). For complete statistics on these scores, see Table 7.

Paired-sample t-tests were run to compare the evaluation performance means (percentage correct) of touchscreen and embossed representations. The analysis revealed no statistically significant differences between the percent means of embossed and touchscreen bar charts ($t(18) = 1.68$, $p = .11$, 95% CI $= -1.32 - 11.85$, Cohen's $d = 0.269$), nor geometric images ($t(21) = -.366$, $p = .718$, 95% CI $= -8.14 - -5.71$, Cohen's $d = 0.073$). These results show that performance when using the multimodal graphics as informed by the guidelines put forth in this article is within the statistical margin of error observed with traditional, embossed tactile graphics.

As the distribution of graphical media, especially in education, moves further into the digital space, this is a promising finding. These guidelines demonstrate that consumption of non-visual media can occur in the digital space, with graphical literacy performance that is at least on par with traditional tactile media. This is an advantageous finding for students in classrooms consuming subject-related material, but even has implication for adults who must read and interpret visualized data, such as bar-graph renderings.

Participants performed better on geometric figure evaluations (embossed: $M = 85.53\%$; multimodal: $M = 86.44\%$) than on the bar-chart evaluations (embossed: $M = 81.82\%$; multimodal: $M = 76.56\%$). A paired-sample t-test was run to determine if the difference in performance between the bar chart and geometric figures was significant. The difference was found to be statistically significant ($t(37) = -2.99$, $p = .005$, 95% CI $= -13.07\% - -2.511\%$, Cohen's $d = 0.430$). This finding is interesting, as it implies that representations of geometric figures, no matter the medium, may

Table 8. Sample Graphic Evaluation Questions

Breakfast Bar Chart	
Multimodal	Embossed
1. How many bars are on the page?	1. How many bars are on the page?
2. What is the title of the graph?	2. What is the title of the graph?
3. What is the x-axis labeled as?	3. What is the y-axis labeled as?
4. How many boys selected grain products?	4. How many girls selected grain products?
5. How many girls selected milk products?	5. How many boys selected milk products?
6. What group selected more milk products?	6. What group selected more milk products?

Quadrilateral Figure	
Multimodal	Embossed
1. How many sides does the figure have?	1. How many sides does the figure have?
2. Are any sides parallel?	2. Are any sides parallel?
3. How did you determine this?	3. How did you determine this?
4. What is the length of side WX?	4. What is the length of side WZ?
5. What is the length of size XY?	5. What is the length of side XY?
6. What is the measure of angle WZY?	6. What is the measure of angle ZYX?

have been easier for our users to understand. Perhaps the complexity of the bar charts contributed to this finding. With the geometric figures, it was a relatively straightforward task to extract and locate information in which to answer questions. Bar charts, however, required users to perform more complex actions, such as locate multiple bars, find bar heights, explore two axes, and evaluate gridlines.

Geometric figure representations, while simpler and more successfully evaluated, were frustrating to the participants. Many remarked it was difficult to create a mental image of the shape due to the amount of information that was required to have in memory throughout the exploration. Whereas bar charts have a distinct set of parts, geometric figures can vary widely in representation. To have a successful representation of the shape, all key parts (all vertices, lines, and special angles) must have been found and stitched together in memory. To assist future users in successfully understanding multimodal, geometric figures, participants had the following suggestions:

(1) White space must be designated important (inside the shape) or unimportant (outside the shape). This is represented in our current iteration with a clicking noise inside the shape, as described in Section 3.
(2) Audio description of angles must be concise and able to be turned off easily should the user have no use for remaining auditory information.
(3) Provide high-level legend of components before exploration, such as a list of all angles, line segments, and any explicit measurements.

Despite having slightly worse performance of the two graphics, many users remarked that the multimodal bar-chart representations were the strongest and easiest to understand. They appreciated the immediacy in which they were able to obtain basic information about the chart. A common suggestion by the participants to make the bar charts easier and quicker was to provide some sort of legend or key for the bar charts, similar to the legend in the tactile version of the chart, and similar to the suggestion for geometric figures above. This would support user search efforts by

Fig. 8. Finger position data of each of the six multimodal images. Bar charts from left to right: Maine Voting Results, Breakfast Results, and Fruits and Vegetables (training). Geometry figures from left to right: right triangle, trapezoid, triangle (training).

allowing them to obtain high-level information about the graphic's data at the start of the task so they might better direct their search of the graphic.

Although direct time comparisons were not measured, experimenters observed that embossed graphics remained quicker to evaluate. This could be due to several reasons with the most obvious being the experience disparity between the two mediums. Participants are likely to have had more experience exploring physical media as compared to the brief training period in which users had to acquaint themselves with the touchscreen-based media. Yet, after gaining familiarity with the bar charts, participants moved very quickly through the touchscreen-based bar charts and users may have eventually caught up to the speeds at which they consumed the embossed media. These observations, however, need to be quantitatively validated.

To demonstrate the most important features of each graphic, heatmaps of participant finger locations are provided by Figure 8. A number of common trends can be seen in these composites.

From the bar-chart renderings, participants spent most of their exploration at the top of the bars, especially at the corners. Participants tended to stay on the cusp of vibratory feedback when exploring large, haptically represented spaces, preferring to trace along the edge between haptics and no haptics. Therefore a large bar space, while visually appealing, may be unnecessary in a non-visual rendering as long as the bar can be readily found and easily traced by users. Very little time was spent at the labels, as participants did not need to explore the entirety of the text strings to obtain immediate audio messages of the labels. This allowed users to easily obtain audio information and quickly move on, enabling them to efficiently use their time.

For rendering of geometric figures, the heatmaps demonstrate a majority of the exploration time spent by participants was to learn the outline of the shape. Participants often returned to key points of interest, anchoring to angles and segments with defined measurements, to make their next exploration decisions. Because of this, not much time was spent tracing the middle sections of long lines—if participants remembered there was a key point nearby, they chose to find the key point from another spatial reference rather than trace the line all the way down. Again, this serves to highlight the importance and the desire for users to use their time efficiently during these tasks.

6.1 Study Limitations

There were limitations to our study. First, we were constrained by time, and thus our within sub-jects trials were low. Additionally, this limited time also abbreviated the amount of training that could be provided to participants. We acknowledge that this study does not capture what may be possible with multimodal graphics over extended training and use. Second, our participant pool was highly varied across age, gender, and relevant graphic skills. While this inserts heterogene-ity within our study, this is a typical sample of the population, although it may not represent all users in each age bracket. Third, the information being evaluated with the graphics was lower-level than many participants had recent experience with or could remember from their primary and secondary school experience. For instance, some participants could not remember the differ-ence between the y- or x-axes nor could some remember how to differentiate between different types of triangles (right, acute, obtuse). Future work will investigate these graphical concepts with school-aged children.

Additionally, after some testing, our team discovered that the text-to-speech rate was too quick for some users. The text-to-speech rate was originally set to 2× the normal human speech rate (as described in the Android documentation). During the experiment, users were given the 2× rate by default, but were provided the option to decrease the rate to 1×. Nonetheless, we note that our text-to-speech rate had limited options for users to choose, which could have affected their time and interpretation during exploration. We also note that it was observed that users pressed quite strongly on the screen when following vibrating lines or bars. This caused the tablet to move slightly on the table during exploration. To mitigate this during the experiment, rubber stickers were placed at the corners of the tablet to prevent the tablet from moving.

Another area that was observed as limiting the study was the double-tap gesture. Several users found the double-tap gesture difficult to trigger. Prior to the study, the time frame for the double-tap was widened to account for slow tappers after in-house testing with local VI individuals suggested this may be necessary. However, in this study, speed was not the issue encountered, but rather positioning. It was observed that users would sometimes not tap twice on the exact same location. When this happens, it is difficult for AndroidOS to register finger touchdown events that do not occur within a couple of millimeters of each other as a double-tap. As such, the gesture was not always activated correctly within the study, which caused confusion for participants. In future work, other alternatives for this gesture will be considered.

7 CONCLUSION

This article puts forth a series of design guidelines, exploration strategies, and best practices for rendering multimodal graphics on touchscreens using vibration, sound, and text feedback. The iterative design and software implementation process is highlighted and example use cases are presented with two foundational math concepts: bar charts and geometry figures. A user study with 22 individuals with visual impairment was conducted comparing performance on informa-tion extraction and comparison tasks between the touchscreen media and traditional, hard-copy embossed graphics. No significant difference in performance was found between mediums. This is interesting, as many participants had training and experience with embossed graphics while having only the training epoch on the multimodal graphics in this study. These findings and the spirit of this work as a whole do not suggest that touchscreen renderings should replace hard-copy, embossed graphics in any way. Rather, we view these results as positive in the light of providing digital access that has similar outcomes to print access. These findings also present more questions for the community to consider. This includes how these guidelines and strategies scale across con-cepts and complexity of graphics and investigations on the limitations of this paradigm. Future

work will explore extensions of these guidelines to maps and diagrams and will also uncover how to deal with graphical concepts that are too complex for a single "multimodal" view. This work also sets the stage for investigations into multisensory experiences, like graphics rendering, for all individuals, not just those with visual impairments, and explorations into new applications for this research to have even broader impact.

ACKNOWLEDGMENTS

REFERENCES

[1] Lopez Ada. 2014. *Reach for the Stars: Touch, Look, Listen, Learn*. SAS Institute. Retrieved from https://books.apple.com/us/book/reach-for-stars-touch-look/id763516126.

[2] Richard Adams, Dianne Pawluk, Margaret Fields, and Ryan Clingman. 2015. Multimodal application for the perception of spaces (MAPS). In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*. ACM, 393–394.

[3] AFB. 2017. Key Employment Statistics for People Who Are Blind or Visually Impaired. Retrieved from https://www.afb.org/research-and-initiatives/statistics/key-employment-statistics.

[4] APH. 2018. Graphiti. Retrieved from http://www.aph.org/graphiti/.

[5] Apple. 2019. Vision Accessibility. Retrieved from https://www.apple.com/accessibility/mac/vision/.

[6] Shiri Azenkot, Cynthia L. Bennett, and Richard E. Ladner. 2013. DigiTaps: Eyes-free number entry on touchscreens with minimal audio feedback. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST'13)*. ACM, New York, NY, 85–90. DOI:https://doi.org/10.1145/2501988.2502056

[7] Catherine M. Baker, Lauren R. Milne, Ryan Drapeau, Jeffrey Scofield, Cynthia L. Bennett, and Richard E. Ladner. 2016. Tactile graphics with a voice. *ACM Trans. Access. Comput.* 8, 1 (2016), 3.

[8] BANA. 2019. Code Books and Guidelines. Retrieved from http://www.brailleauthority.org/publications-area.html.

[9] BLC. 2019. Braille Literacy Canada. Retrieved from http://www.brailleliteracycanada.ca/en/.

[10] Stephen A. Brewster and Lorna M. Brown. 2004. Tactons: Structured tactile messages for non-visual information display. In *Proceedings of the 5th Conference on Australasian User Interface-Volume 28*. Australian Computer Society, Inc., 15–23.

[11] Stephen A. Brewster, Faraz Chohan, and Lorna Brown. 2007. Tactile feedback for mobile interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 159–162.

[12] 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*. ACM, 107–114.

[13] Maria Claudia Buzzi, Marina Buzzi, Barbara Leporini, and Amaury Trujillo. 2017. Analyzing visually impaired people's touch gestures on smartphones. *Multimedia Tools Applic.* 76, 4 (2017), 5141–5169.

[14] Seungmoon Choi and Katherine J. Kuchenbecker. 2013. Vibrotactile display: Perception, technology, and applications. *Proc. IEEE* 101, 9 (2013), 2093–2104.

[15] Rafael Jeferson Pezzuto Damaceno, Juliana Cristina Braga, and Jesús Pascual Mena-Chalco. 2018. Mobile device accessibility for the visually impaired: Problems mapping and recommendations. *Univ. Access Inf. Soc.* 17, 2 (2018), 421–435.

[16] Core Haptics | Apple Developer. 2020. Retrieved from https://developer.apple.com/documentation/corehaptics.

[17] Julie Ducasse, Anke M. Brock, and Christophe Jouffrais. 2018. Accessible interactive maps for visually impaired users. In *Mobility of Visually Impaired People*. Springer, 537–584.

[18] Christin Engel and Gerhard Weber. 2018. A user study to evaluate tactile charts with blind and visually impaired people. In *Proceedings of the International Conference on Computers Helping People with Special Needs*. Springer, 177–184.

[19] William Erickson, Camille Lee, and Sarah von Schrader. 2012. Disability Statistics from the 2011 American Community Survey (ACS). Retrieved from www.disabilitystatistics.org.

[20] George A. Gescheider, John H. Wright, and Ronald T. Verrillo. 2010. *Information-processing Channels in the Tactile Sensory System: A Psychophysical and Physiological Analysis*. Psychology Press.

[21] Frédéric Giraud, Tomohiro Hara, Christophe Giraud-Audine, Michel Amberg, Betty Lemaire-Semail, and Masaya Takasaki. 2018. Evaluation of a friction reduction based haptic surface at high frequency. In *Proceedings of the IEEE Haptics Symposium (HAPTICS'18)*. IEEE, 210–215.

[22] Nicholas A. Giudice, Benjamin A. Guenther, Nicholas A. Jensen, and Kaitlyn N. Haase. 2020. Cognitive mapping without vision: Comparing wayfinding performance after learning from digital touchscreen-based multimodal maps vs. embossed tactile overlays. *Front. Hum. Neurosci.* 14 (2020), 87.

[23] Nicholas A. Giudice, Hari Prasath Palani, Eric Brenner, and Kevin M. Kramer. 2012. Learning non-visual graphical information using a touch-based vibro-audio interface. In *Proceedings of the 14th International ACM SIGACCESS Conference on Computers & Accessibility*. ACM, 103–110.

[24] Cagatay Goncu and Kim Marriott. 2011. GraVVITAS: Generic multi-touch presentation of accessible graphics. In *Proceedings of the Conference on Human-Computer Interaction*. Springer, 30–48.

[25] Jenna L. Gorlewicz, Jessica Burgner, Thomas J. Withrow, and Robert J. Webster III. 2014. Initial experiences using vibratory touchscreens to display graphical math concepts to students with visual impairments. *J. Special Educ. Technol.* 29, 2 (2014), 17–25.

[26] Jenna L. Gorlewicz, Jennifer L. Tennison, Hari Prasath Palani, and Nicholas A. Giudice. 2018. The graphical access challenge for people with visual impairments: Positions and pathways forward. In *Interactive Multimedia*. IntechOpen.

[27] W. Grussenmeyer and E. Folmer. 2017. Accessible touchscreen technology for people with visual impairments: A survey. *Trans. Access. Comput.* 9, 2 (2017), 6:1–6:31.

[28] João Guerreiro and Daniel Gonçalves. 2015. Faster text-to-speeches: Enhancing blind people's information scanning with faster concurrent speech. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*. 3–11.

[29] Tiago Guerreiro, Kyle Montague, João Guerreiro, Rafael Nunes, Hugo Nicolau, and Daniel J. V. Gonçalves. 2015. Blind people interacting with large touch surfaces: Strategies for one-handed and two-handed exploration. In *Proceedings of the International Conference on Interactive Tabletops & Surfaces*. 25–34.

[30] Michael Hahn, Corrine Mueller, and Jenna L. Gorlewicz. 2019. The comprehension of stem graphics via a multisensory tablet in students with visual impairment. *J. Vis. Impair. Blind.* 113, 5 (2019).

[31] Yvette Hatwell, Arlette Streri, and Edouard Gentaz. 2003. *Touching for Knowing: Cognitive Psychology of Haptic Manual Perception.* Vol. 53. John Benjamins Publishing.

[32] Eve Hoggan, Stephen A. Brewster, and Jody Johnston. 2008. Investigating the effectiveness of tactile feedback for mobile touchscreens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1573–1582.

[33] Inwook Hwang and Seungmoon Choi. 2010. Perceptual space and adjective rating of sinusoidal vibrations perceived via mobile device. In *Proceedings of the IEEE Haptics Symposium*. IEEE, 1–8.

[34] Apple iOS. 2019. Haptics—User Interaction. Retrieved from https://developer.apple.com/design/human-interface-guidelines/ios/user-interaction/haptics/.

[35] Yvonne Jansen, Thorsten Karrer, and Jan Borchers. 2010. MudPad: Tactile feedback and haptic texture overlay for touch surfaces. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*. ACM, 11–14.

[36] Lynette A. Jones and Susan J. Lederman. 2006. *Human Hand Function.* Oxford University Press.

[37] Shaun K. Kane, Jeffrey P. Bigham, and Jacob O. Wobbrock. 2008. Slide rule: Making mobile touch screens accessible to blind people using multi-touch interaction techniques. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, 73–80.

[38] Shaun K. Kane, Meredith Ringel Morris, Annuska Z. Perkins, Daniel Wigdor, Richard E. Ladner, and Jacob O. Wobbrock. 2011. Access overlays: Improving non-visual access to large touch screens for blind users. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST'11)*. ACM, New York, NY, 273–282. DOI:https://doi.org/10.1145/2047196.2047232

[39] Roberta L. Klatzky, Nicholas A. Giudice, Christopher R. Bennett, and Jack M. Loomis. 2014. Touch-screen technology for the dynamic display of 2D spatial information without vision: Promise and progress.*Multisens. Res.* 27 (2014), 359–378.

[40] Roberta L. Klatzky and Susan J. Lederman. 2003. Touch. In *Handbook Psychol.* A. F. Healy and R. W. Proctor (Eds.). John Wiley & Sons, 147–176.

[41] Susan J. Lederman and Roberta L. Klatzky. 1987. Hand movements: A window into haptic object recognition. *Cog. Psychol.* 19 (1987), 342–368.

[42] Susan J. Lederman and Roberta L. Klatzky. 2009. Haptic perception: A tutorial. *Atten. Percept. Psychophys.* 71, 7 (2009), 1439–1459.

[43] Jack M. Loomis, Roberta L. Klatzky, and Nicholas A. Giudice. 2018. Sensory substitution of vision: Importance of perceptual and cognitive processing. In *Assistive Technology for Blindness and Low Vision*. CRC Press, 179–210.

[44] Jack M. Loomis and Susan J. Lederman. 1986. Tactual perception. *Handbook Percept. Hum. Perform.* 2 (1986), 2.

[45] James Minogue and M. Gail Jones. 2006. Haptics in education: Exploring an untapped sensory modality. *Rev. Educ. Res.* 76, 3 (2006), 317–348.

[46] Annika Muehlbradt, Madhur Atreya, Darren Guinness, and Shaun K. Kane. 2018. Exploring the design of audio-kinetic graphics for education. In *Proceedings of the International Conference on Multimodal Interaction.* ACM, 455–463.

[47] Joseph Mullenbach, Michael A. Peshkin, and J. Edward Colgate. 2016. eShiver: Lateral force feedback on fingertips through oscillatory motion of an electroadhesive surface. *IEEE Trans. Haptics* 10, 3 (2016), 358–370.

[48] Michael A. Nees and Bruce N. Walker. 2008. Encoding of information in auditory displays: Initial research on flexibility and processing codes in dual-task scenarios. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 52. SAGE Publications, Los Angeles, CA, 1820–1824.

[49] Michael A. Nees and Bruce N. Walker. 2009. Auditory Interfaces and Sonification. 1–15. http://sonify.psych.gatech.edu/~walkerb/publications/pdfs/NeesWalker-UniversalAccessChapter2007-submitted.pdf.

[50] Uran Oh, Stacy Branham, Leah Findlater, and Shaun K. Kane. 2015. Audio-based feedback techniques for teaching touchscreen gestures. *ACM Trans. Access. Comput.* 7, 3 (2015), 9.

[51] Sile O'Modhrain, Nicholas A. Giudice, John A. Gardner, and Gordon E. Legge. 2015. Designing media for visually-impaired users of refreshable touch display: Possibilities and pitfalls. *Trans. Haptics* 8, 3 (2015), 248–257.

[52] Hari Prasath Palani and Nicholas A. Giudice. 2014. Evaluation of non-visual panning operations using touch-screen devices. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility.* ACM, 293–294.

[53] Hari Prasath Palani and Nicholas A. Giudice. 2017. Principles for designing large-format refreshable haptic graphics using touchscreen devices: An evaluation of nonvisual panning methods. *ACM Trans. Access. Comput.* 9, 3 (2017), 9.

[54] Hari Prasath Palani, Uro Giudice, and Nicholas A. Giudice. 2016. Evaluation of non-visual zooming operations on touchscreen devices. In *Proceedings of the International Conference on Universal Access in Human-Computer Interaction.* Springer, 162–174.

[55] Hari Prasath Palani, Jennifer L. Tennison, G. Bernard Giudice, and Nicholas A. Giudice. 2018. Touchscreen-based haptic information access for assisting blind and visually-impaired users: Perceptual parameters and design guidelines. In *Proceedings of the International Conference on Applied Human Factors and Ergonomics.* Springer, 837–847.

[56] Pew. 2019. Demographics of Mobile Device Ownership and Adoption in the United States. Retrieved from https://www.pewinternet.org/fact-sheet/mobile/.

[57] View Plus. 2018. View Plus. Retrieved from http://www.viewplus.com.

[58] Edman Polly. 1992. *Tactile Graphics.* American Foundation for the Blind. New York.

[59] Benjamin Poppinga, Charlotte Magnusson, Martin Pielot, and Kirsten Rassmus-Gröhn. 2011. TouchOver map: Audio-tactile exploration of interactive maps. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services.* ACM, 545–550.

[60] Monoj Kumar Raja. 2011. *The Development and Validation of a New Smartphone Based Non-visual Spatial Interface for Learning Indoor Layouts.* Master's Thesis. The University of Maine.

[61] Ravi Rastogi and Dianne T. V. Pawluk. 2010. Automatic, intuitive zooming for people who are blind or visually impaired. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility.* ACM, 239–240.

[62] Ravi Rastogi and Dianne T. V. Pawluk. 2013. Toward an improved haptic zooming algorithm for graphical information accessed by individuals who are blind and visually impaired. *Assist. Technol.* 25, 1 (2013), 9–15.

[63] Ravi Rastogi, T. V. Dianne Pawluk, and Jessica Ketchum. 2013. Intuitive tactile zooming for graphics accessed by individuals who are blind and visually impaired. *IEEE Trans. Neural Syst. Rehab. Eng.* 21, 4 (2013), 655–663.

[64] Alexander Russomanno, R. Brent Gillespie, Sile O'Modhrain, and Mark Burns. 2015. The design of pressure-controlled valves for a refreshable tactile display. In *Proceedings of the IEEE World Haptics Conference (WHC'15).* IEEE, 177–182.

[65] Freedom Scientific. 2019. Freedom Scientific. Retrieved from https://www.freedomscientific.com/.

[66] Joseph C. Stevens, Emerson Foulke, and Matthew Q. Patterson. 1996. Tactile acuity, aging, and Braille reading in long-term blindness. *J. Exper. Psychol.: Appl.* 2, 2 (1996), 91.

[67] Ed Summers, Julianna Langston, Robert Allison, and Jennifer Cowley. 2012. Using SAS/GRAPH to create visualizations that also support tactile and auditory interaction. In *Proceedings of the SAS Global Forum.* Citeseer.

[68] Hong Z. Tan, Nathaniel I. Durlach, Charlotte M. Reed, and William M. Rabinowitz. 1999. Information transmission with a multifinger tactual display. *Percept. Psychophys.* 61, 6 (1999), 993–1008.

[69] Tanvas. 2019. Tanvas. Retrieved from https://tanvas.co/.

[70] Blitab Technology. 2018. Blitab. Retrieved from http://blitab.com/.

[71] Jennifer L. Tennison, Zachary S. Carril, Nicholas A. Giudice, and Jenna L. Gorlewicz. 2018. Comparing haptic pattern matching on tablets and phones: Large screens are not necessarily better. *Optom. Vis. Sci.* 95, 9 (2018), 720–726.

[72] Jennifer L. Tennison and Jenna L. Gorlewicz. 2016. Toward non-visual graphics representations on vibratory touch-screens: Shape exploration and identification. In *Proceedings of the International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 384–395.

[73] Jennifer L. Tennison and Jenna L. Gorlewicz. 2019. Non-visual perception of lines on a multimodal touchscreen tablet. *ACM Trans. Appl. Percept.* 16, 1 (2019), 6.

[74] Jennifer L. Tennison, P. Merlin Uesbeck, Nicholas A. Giudice, Andreas Stefik, Derrick W. Smith, and Jenna L. Gorlewicz. 2020. Establishing vibration-based tactile line profiles for use in multimodal graphics. *Trans. Appl. Percept.* 17, 2 (2020), 1–14.

[75] American Thermoform. 2019. Swell Touch Paper. Retrieved from http://www.americanthermoform.com/product/swell-touch-paper/.

[76] Robert W. Van Boven, Roy H. Hamilton, Thomas Kauffman, Julian P. Keenan, and Alvaro Pascual-Leone. 2000. Tactile spatial resolution in blind Braille readers. *Neurology* 54, 12 (2000), 2230–2236.

[77] Yasemin Vardar, Burak Güçlü, and Cagatay Basdogan. 2017. Effect of waveform on tactile perception by electrovibration displayed on touch screens. *IEEE Trans. Haptics* 10, 4 (2017), 488–499.

[78] ViTAL. 2019. Retrieved from https://www.vital.education/.

[79] Bruce N. Walker. 2002. Magnitude estimation of conceptual data dimensions for use in sonification. *J. Exper. Psychol.: Appl.* 8, 4 (2002), 211.

[80] Bruce N. Walker, Jeffrey Lindsay, Amanda Nance, Yoko Nakano, Dianne K. Palladino, Tilman Dingler, and Myounghoon Jeon. 2013. Spearcons (speech-based earcons) improve navigation performance in advanced auditory menus. *Hum. Fact.* 55, 1 (2013), 157–182.

[81] Bruce N. Walker and Lisa M. Mauney. 2010. Universal design of auditory graphs: A comparison of sonification mappings for visually impaired and sighted listeners. *ACM Trans. Access. Comput.* 2, 3 (2010), 12.

[82] WHO. 2011. Visual Impairment and Blindness. Fact Sheet No. 282. Retrieved from http://www.who.int/mediacentre/factsheets/fs282/en/.

[83] Maarten W. A. Wijntjes, Thijs Van Lienen, Ilse M. Verstijnen, and Astrid M. L. Kappers. 2008. Look what I have felt: Unidentified haptic line drawings are identified after sketching. *Acta Psycholog.* 128, 2 (2008), 255–263.

[84] Heng Xu, Michael A. Peshkin, and J. Edward Colgate. 2019. UltraShiver: Lateral force feedback on a bare fingertip via ultrasonic oscillation and electroadhesion. *IEEE Trans. Haptics* 12, 4 (2019), 497–507.

[85] Koji Yatani and Khai Nhut Truong. 2009. SemFeel: A user interface with semantic tactile feedback for mobile touch-screen devices. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology*. ACM, 111–120.

[86] Wai Yu and Stephen A. Brewster. 2002. Comparing two haptic interfaces for multimodal graph rendering. In *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'02)*. IEEE, 3–9.

[87] Wai Yu and Stephen A. Brewster. 2003. Evaluation of multimodal graphs for blind people. *Univ. Access Inf. Soc.* 2, 2 (2003), 105–124.