RoboGraphics: Dynamic Tactile Graphics Powered by Mobile Robots

Darren Guinness, Annika Muehlbradt, Daniel Szafir, Shaun K. Kane

University of Colorado Boulder

Boulder, CO USA

{darren.guinness, annika.muehlbradt, daniel.szafir, shaun.kane}@colorado.edu

ABSTRACT

Tactile graphics are a common way to present information to people with vision impairments. Tactile graphics can be used to explore a broad range of static visual content but aren't well suited to representing animation or interactivity. We introduce a new approach to creating dynamic tactile graphics that combines a touch screen tablet, static tactile overlays, and small mobile robots. We introduce a prototype system called *RoboGraphics* and several proof-of-concept applications. We evaluated our prototype with seven participants with varying levels of vision, comparing the RoboGraphics approach to a flat screen, audio-tactile interface. Our results show that dynamic tactile graphics can help visually impaired participants explore data quickly and accurately.

Author Keywords

Accessibility; Tangible User Interfaces; Education; Tactile; Robots; Blindness.

CCS Concepts

Human-centered computing → Accessibility systems and tools.

INTRODUCTION

Being able to read and understand charts, diagrams, and other graphics is an important life skill. In recent years, simple charts have often been replaced by more complex information graphics, animations, and interactive data visualizations. Each of these media types can present accessibility challenges to blind or visually impaired people. For simple images and graphics, it may be possible to create an accessible version of that media using sonification [[23]], audio and text descriptions [[27]], or tactile graphics [[3],[4]].

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

ASSETS '19, October 28–30, 2019, Pittsburgh, PA, USA

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

ACM 978-1-4503-6676-2/19/10...\$15.00

<https://doi.org/10.1145/3308561.3353804>

Figure 1. In the Braille Assistant app, mobile robots form a large, tactile version of the letter “Q” (⠠) in Braille.

In many cases, the translation from visual content into a non-visual format is made by a teacher or Braille transcriptionist, often with limited time and limited access to technology. By necessity, these translations are often simple and lack many details present in the source material. For example, tactile graphics made by educators are often constructed from “low-tech” materials such as paper or Lego bricks. These tools often present information at a low resolution, and do not support advanced features such as animation or interactivity.

While most tactile graphics “in the wild” rely upon simple static materials, researchers have developed many more advanced technologies for representing information non-visually, including talking tactile graphics [[4],[18]] and interactive shape displays [[6],[7],[10],[20]]. However, these systems often rely on technologies that are expensive and difficult to use, if not completely unavailable to consumers. This gap between low-tech, practical materials and high-tech, difficult to access prototypes, suggests a need for tools to enable the creation of more complex non-visual graphics using readily available technology.

In this paper, we explore a new approach to creating dynamic tactile graphics using mobile robots. This approach, called *RoboGraphics*, combines a touch screen tablet, a static tactile overlay, and small mobile robots to implement a shape-changing display using low-cost and off-the-shelf technology. Our software platform supports tactile graphics that combine static and moving tactile content with audio feedback. These dynamic tactile graphics can be controlled

by the user via touch and speech input. We demonstrate the breadth of this approach through several proof-of-concept applications including animated graphics, interactive tutors, and audio-tactile stories.

We evaluated our prototype with 7 participants with vision impairments to investigate two questions. First, how do dynamic tactile graphics compare to a flat touch screen for exploring data? Second, how can this approach be extended to support various types of audio-visual media? The results of our study show that dynamic tactile graphics are preferred to a flat screen with audio feedback alone. In addition, our study participants suggested a variety of uses for this technology, including data graphics and interactive stories.

RELATED WORK

Tactile Graphics

Tactile graphics can be useful in presenting spatial information in an accessible format. Tactile graphics can be used to represent maps [[4], [5]], data visualizations [[3]], Braille characters [[2]], interactive stories [[19]], and other visual media. Some research prototypes have explored adding audio descriptions to tactile materials (e.g., [[4],[18]]). Two major limitations of most tactile graphics are that they cannot easily represent moving images such as animations or video, and that they usually do not support interaction from the user.

Add-On Haptic Devices

One way to add additional haptic feedback to a traditional interface is via an add-on haptic device, such as a handheld puck [[16]], a haptic stylus [[25]], or a mobile robot [[9]]. With this approach, the user holds the device in their hand while exploring a 2D surface or 3D space. This approach can be augmented with audio feedback to provide a multimodal experience [[25],[26]]. While these devices enable the addition of haptic feedback without incorporating a high-cost shape display, they are usually limited to one or two points of feedback and may be limited in the types of feedback that they can provide, such as only providing vibration feedback.

Shape Displays

Shape displays, which can represent an image in tactile format, are an increasingly common solution to visual accessibility problems. Most current shape displays use a grid of pins that can individually raise or lower to form complex shapes [[1],[7],[10]]. However, some shape displays provide mobile elements, such as a tactile markers, that can propel themselves over the display surface [[8],[14],[15],[17],[18],[20]]. These displays often have the ability to convey motion by allowing a user to touch the tactile marker as it moves. Currently, most shape displays of this type use custom hardware which can be difficult to obtain or costly [[1],[6],[7],[17],[20]]. Our approach uses off-the-shelf components that can be purchased and assembled by end users.

The RoboGraphics system builds off of the earlier Haptic Video Player system [8]. However, while the Haptic Video

Player used mobile robots to represent video, RoboGraphics extends this approach to support interactive applications, and supports interactions between static and moving tactile components.

ROBOGRAPHICS

The *RoboGraphics* system (Figure 1) comprises the following components: a touch screen tablet, a set of reusable mobile robots, and a series of static tactile overlays. When displaying a graphic, animation, or application, the user places a tactile overlay on the screen, and then places one or more of the mobile robots onto the overlay.

These three components work together to create an interactive tactile graphic: the touch screen accepts user input and displays control codes for the mobile robots, the static tactile overlay represents static portions of the graphic, and the robots move across the screen to represent moving and interactive elements.

Hardware Components

The RoboGraphics prototype uses a combination of off-the-shelf electronics (tablet, mobile robots) and low-cost tactile materials (tactile overlays and tactile mounts).

Touch Screen Tablet

The current prototype uses a Lenovo Yoga 720 convertible laptop, with a 15.6-inch, 16:9 touch screen, running Windows 10. However, this system can be modified to work on any touch screen-based laptop or tablet device.

Mobile Robots

Dynamic tactile feedback is provided by a set of Ozobot Bit 2.0 robots (Figure 2). These robots are approximately one cubic inch in volume. The robots receive commands via an optical sensor on the bottom and drive using two motorized wheels. We chose these robots because of their programmability, ergonomic size, and low cost (about \$60 USD each).

Our proof-of-concept applications use between one and six robots each. To keep the cost of entry low, we designed the RoboGraphics system so that it could be useful with only a single robot. Adding more robots enables the system to present multiple moving objects and to represent more complex shapes, such as Braille characters.

Figure 2. Mobile robots used in the RoboGraphics system. The Ozobot Bit 2.0 is driven by two wheels and reads commands from the optical sensor on its bottom.

Tactile Overlays

While a major benefit of RoboGraphics is its ability to represent moving data via the mobile robots, not all tactile information is best represented using a robot. When designing RoboGraphics, we found many situations in which information was most useful if it remained still throughout the interaction, such as the tick marks on a bar chart. Furthermore, the mobile robot hardware is not able to represent some fine details such as contours or textures. To address this limitation, RoboGraphics combines mobile robots with a static tactile overlay component.

Currently, each tactile overlay consists of a single layer of material with cutouts and etchings to indicate tactile features (Figure 1). To properly align the overlays and keep them still during interactions with the system, we created a static frame that we attached to the touch screen using tape. Each overlay features a notched corner in the lower-right which helps the user to orient the overlay properly. We created the initial set of overlays by laser cutting cardboard; however, these overlays could be made from other materials and could be cut by hand if a laser cutter (or equivalent tool) is not available.

In general, the tactile overlay is meant to represent the background and context of the dynamic tactile graphic, including points of reference such as graph axes, legends, decorative background graphics, and other tactile information. We used stick-on Braille labels to add textual information to the overlays.

In addition to providing tactile information, we leverage the tactile overlays to support non-visual interactions with RoboGraphics. For example, each overlay includes tactile “robot bays,” cutouts that are the approximate shape and size of the mobile robots and which indicate the robot’s starting position. When loading a new graphic, the user can feel the tactile cues and place the robot in the appropriate location. RoboGraphics also supports interactivity via the addition of tactile “button” cutouts, in which the user presses their finger into a circular cutout on the overlay to perform some function.

Tactile Mounts

An additional limitation of the mobile robot form factor is that all robots share the same size and shape (a hemisphere approximately one cubic inch in volume). In prior research [[8]], we found that it was sometimes difficult to track multiple robots simultaneously and to remember which robot represented which object.

To address this problem, RoboGraphics includes support for attaching static “tactile mounts” to the top of each robot. These tactile mounts give each robot a distinguishing shape that enables the user to more easily recognize the robot and determine its function (Figure 3). For example, for the Tortoise and the Hare storybook application, one robot is given the silhouette of a turtle while the other is given the silhouette of a rabbit.

As with the tactile overlays, we created the tactile mounts by laser cutting cardboard sheets. We attached the mounts to the tops of the robots using velcro. Once again, it would be possible to develop tactile mounts using other materials and fabrication methods, or even to create tactile mounts with more elaborate three-dimensional shapes.

Figure 3. Tactile mounts may be placed on the top of robots to provide distinguishing shapes and textures. These mounts represent the storybook characters of the Tortoise and the Hare.

System Software

RoboGraphics are controlled by a custom C# application. This application is responsible for controlling the motion of the robots, playing speech and other audio, and accepting user input.

Currently, each “application” is made up of additional C# code that is integrated with the main program. Developing user-friendly authoring tools for RoboGraphics is an important goal of our future work.

Robot Control

Each robot is controlled via an optical sensor placed at its bottom. By default, each robot follows lines underneath it until it reaches the end of the line or a color code. These color codes are part of the Ozobot hardware platform, and include commands for adjusting speed, performing a U-turn, and twirling (Table 1).

While the color codes themselves are part of the Ozobot platform, the RoboGraphics software provides an interface for a programmer to specify the location each robot should travel to, and dynamically draws paths on the screen to move the robot from its starting point to the intended end point. Thus, when creating an audio-tactile story, the programmer can direct the robots to a specific location, read out some text, and then move the robot to a new location.

One limitation of this approach is that the system does not track the actual positions of the robots, as there is no overhead camera or other sensor to track their location. Thus, if a user knocks a robot out of position, the user will have to notice this problem and correct the problem themselves. In RoboGraphics, this typically involves placing the robot back

in its starting position (indicated by the tactile overlay) and pressing a key to restart the current animation.

User Interaction

In addition to moving the robots across the screen, RoboGraphics provides output via text-to-speech and prerecorded sound files.

Command	Color Code
Slow (25mm/sec)
Cruise (45mm/sec)
Fast (65 mm/sec)
Turbo (85 mm/sec)
U-turn
Twirl

Table 1. Graphical commands sent to the robots via the underlying touch screen.

User input can be provided by touching portions of the screen marked as virtual buttons. To create a button with tactile feedback, we create a circular cutout in the tactile overlay, add a Braille label to the overlay, and detect when the user touches in that area. To reduce the risk of accidentally activating a button, the user must touch and hold the screen for 1 second to activate the virtual button. In some applications the user may also input brief voice commands.

PROOF-OF-CONCEPT APPLICATIONS

Our initial application set demonstrates the types of interactions supported by the RoboGraphics system, including static tactile graphics, dynamic motion, and user input, as well as demonstrating how this approach can be used to create a variety of application experiences.

When brainstorming ideas for applications, we considered several types of interactions that would work especially well for this form factor:

Data graphics. The system can be used to represent charts and graphs. The static tactile overlay can be used to provide a frame of reference, while the mobile robots can reconfigure themselves to show a particular dataset.

Audio-tactile stories. The system can be used to play back some narrative content such as a storybook or educational animation. The mobile robots can be used to represent characters in the story or can direct the user's attention to a specific location. The static tactile overlay can be used to provide scenery or other contextual information.

Interactive applications. A set of robots can be used to represent particular shapes, such as letters, numbers, or chemical diagrams. The user can browse through a set of

items to learn about each one and can interact with the system to complete a quiz (e.g., “identify the letter”) or other activity.

We developed an initial set of five applications, shown in Figure 4 and described in detail below.

Tactile Bar Chart

The Tactile Bar Chart application allows a user to explore a series of bar charts that represent some dataset. The user can page between different charts to compare them.

Static Tactile Overlay

The overlay shows four columns to represent data. Tactile lines along the edge of the chart represent the chart's tick marks and can be used to understand the scale of the chart. A row of button cutouts is placed along the left edge of the overlay.

Dynamic Tactile Elements

This application uses four robots. When starting the application, the system prompts the user to place one robot at each of the tactile starting locations. When loading a chart, the robots move to represent the top of each bar in the chart. The user can touch each of the robots to gain an overall sense of the data trend.

User Interaction

After loading the application, the system reads out a brief description of the data including the chart title, the maximum and minimum scale values, and the interval between each tactile marker.

The user can explore the chart by touching the screen. Touching within a column reads out that column's label and current value (value not available during the experiment). The user can press buttons along the left edge to move between different datasets or views.

Tortoise and the Hare

The Tortoise and the Hare is a popular fable that was likely written by Aesop around 550 BCE [[11]]. The story concerns a race between two characters: the slow but determined Tortoise and the fast but lazy Hare. In the story, the Tortoise challenges the Hare to a race. The Hare laughs at the Tortoise's proposition but agrees to participate. When the race begins, the Hare quickly moves far ahead of his opponent. Feeling confident about his success, the Hare then decides to eat some carrots and take a nap under a tree. While the Hare rests, the Tortoise moves slowly and steadily forward, eventually passing the sleeping Hare. When the Hare awakens, he sees that the Tortoise has almost made it to the finish line. The Hare rushes ahead, but the Tortoise maintains the lead and wins the race. This application presents the story of the Tortoise and the Hare via sound and tactile movement.

Static Tactile Overlay

The tactile overlay presents a race track as depicted in the story. Along the top of the track are several tactile reliefs that represent different landmarks along the path. The first two

Figure 4. Applications demonstrated in the study. From left to right: (1) Tactile Bar Chart; (2) Tortoise and the Hare; (3) Analog Clock; (4) Braille Assistant; (5) Digestive System.

tactile images, a mountain and a house, are decorative, while a tactile carrot patch and tree represent elements from the story. The tactile overlay provides separate path for the two racers. A flag is etched into the end of the overlay to denote the finish line.

Dynamic Tactile Elements

This application uses two robots, one to represent each character in the story. Each robot is given a tactile mount so that the Tortoise is shaped like a turtle and the Hare is shaped like a rabbit.

During the story, the two characters move along with spoken narration. For example, the Hare moves ahead of the tortoise at the start of the race, but pauses underneath the tree, while the Tortoise moves steadily forward and eventually reaches the finish line.

User Interaction

The user begins by placing the tactile mounts on each robot and placing the robots at the starting position.

In the lower left corner of the overlay, there are three buttons (labeled with Braille) representing *Play*, *Pause*, and *Replay*, respectively. The story begins when the user presses the *Play* button.

Once started, the story is performed through speech and sound effects. The two characters move in correspondence with their actions in the story. The user can touch the characters to track their location.

Analog Clock

Learning to tell time on an analog clock face can be challenging for both blind and sighted children. To support children in learning how to tell time, we developed an application which conveys time on the analog clock face.

Static Tactile Overlay

The overlay presents a circular clock face. Around the clock face, each hour is denoted with a small tactile cutout and a Braille label.

Dynamic Tactile Elements

This application uses two robots. One robot represents the hour hand and the other represents the minute hand. Each robot is augmented by a tactile mount representing a pointed clock hand. At the start of the application, the robot representing the minute hand moves inward toward the center of the circle, making it easier to differentiate the two hands.

User Interaction

Currently, the system enables the user to view different times as represented on the analog clock face. The user speaks a time (e.g., “12:40”) and the robots move to reflect that time. Touching the screen reads back the current time.

Braille Assistant

The Braille Assistant application demonstrates how RoboGraphics can be used as a learning tool, in this case as a tool to learn and practice Braille letters.

Static Tactile Overlay

The tactile overlay provides six tactile indicators for placing the robots, a large area for representing Braille characters, and buttons for navigating between characters.

Dynamic Tactile Elements

This application uses six robots, one to represent each dot in a standard Braille character. When a character is selected, the relevant dots move to the center of the screen, while the other dots remain at the edge. For example, when representing the Braille character “C” (⠠), the two robots on the top row will move toward the center.

User Interaction

The user is able to view each character of the Braille alphabet. Pressing the *Next* and *Previous* buttons moves to the next or previous character, respectively, while pressing the *Random* button selects a random character.

When a character is displayed, the user reaches toward the center of the screen to feel the tactile representation of the character. Tapping the screen reads out the character name.

Digestive System

The final application demonstrates how RoboGraphics can be used to present a multimodal educational experience. This application presents an interactive exploration of a cow’s digestive system, enabling the user to observe how it eats, digests, and eliminates its food.

This application was inspired by similar demonstration from TinyBop’s educational iOS apps [22] and the Magic School Bus television show [[22]]. While these prior examples demonstrated the digestive process visually, we use RoboGraphics to present this information in an accessible and multimodal format.

Static Tactile Overlay

The tactile overlay shows a cutout silhouette of a cow. The cow’s body is hollow, as the food will pass through this area.

Figure 5. Bar charts presented in the study. From left to right: (1) Student distribution between courses; (2) Number of tickets sold at a movie theatre; (3) Number of vehicles built in a factory; (4) Eye color count in a classroom.

Braille labels indicate the start of the digestive process (near the cow's mouth) and the end of the process (near the cow's rear). A *Next* button is placed in the lower right corner of the screen.

Dynamic Tactile Elements

This application uses a single robot. The robot represents the cow's food as it moves through the digestive system.

User Interaction

This application is presented as an interactive slideshow as the food is chewed, digested, re-digested, passed through the intestines, and eliminated. Pressing the *Next* button moves to the next step. At each step, the robot moves to the next location. To represent this process in a multimodal format, the system sonifies this process through a series of squeaks, squishes, and gurgling sounds, and then provides a spoken word description of what happens at that stage. Tapping within the cow's body reads out a description of nearby organs.

USER STUDY

We conducted a user study to evaluate the RoboGraphics prototype and proof-of-concept applications. Our study was motivated by two questions:

- 1) Can visually impaired people correctly interpret data presented via RoboGraphics?
- 2) What applications of RoboGraphics would be most useful and interesting to end users?

We explored these two questions through a single-session, two-part study.

Participants

We recruited 7 participants (4 male, 3 female, age 29–49) via university mailing lists and via local blindness organizations. Participants' level of vision varied: P1 described his vision as "no usable vision"; P2, P3, and P6 reported their vision as "no vision" or "totally blind" while P4, P5, and P7 reported "some light perception". Participants' occupations included entrepreneur, art teacher, stay-at-home parent, accessibility architect, accessibility consultant, and unemployed.

Procedure

Each study session took between 65 and 90 minutes. After completing the consent process, the researcher presented the tablet and the robots to the participant, allowing the

participant to touch the objects and become familiar with them. The remainder of the study comprised two tasks.

Task 1: Chart Comprehension

To assess comprehension of data presented via RoboGraphics, we presented participants with a series of bar charts in the Tactile Bar Chart application. For each chart, we asked participants to explore the chart, and then asked comprehension questions such as the minimum value, maximum value, and general trend of the chart.

To assess the performance of RoboGraphics relative to other approaches, we presented charts both via RoboGraphics and via a control condition consisting of the touch screen and passive tactile overlay, but with no robots. In this condition, the system provided spoken directions toward the nearest data point [12]. When a user placed their finger on the display, the system began to guide the user in the form of audio instructions ("up, up, left, left, ding!") to the nearest data point. In both cases, touching the data point would read out the column label.

During this task, participants first explored a practice dataset with one interface (audio or tactile) and another practice dataset with the other interface. Then, each participant explored the four charts shown in Figure 5. All participants explored the charts in the same order: half explored charts 1 and 2 with audio (and 3 and 4 with tactile), while half used tactile on charts 1 and 2 (and audio on charts 3 and 4). After exploring each chart for as long as they wished, the researcher asked the participant to identify the maximum, minimum, and trend of the chart. After exploring all of the datasets, the researcher asked the participant to rate each of the interface conditions in terms of enjoyability and informativeness.

Task 2: Application Feedback

For the next task, participants interacted with each of the proof-of-concept applications described previously and discussed each application with the researcher. Each participant tested the applications in the same order: Tortoise and the Hare, Analog Clock, Braille Assistant, and Digestive System.

To familiarize the participant with the interface, the researcher described the purpose of the application and presented the participant with a cardboard cutout of the user interface (Figure 6). These cutouts were the same as the

tactile overlays but mounted on a sheet of cardboard instead of a touch screen. The participant was encouraged to familiarize themselves with the application layout before beginning the next step of the study.

After familiarizing themselves with the screen layout, the participant tested the application using the RoboGraphics system. Once the participant felt that they understood the application, they were asked to describe their impressions of that application before moving on to the next one. After testing all of the applications, the participant provided general feedback about their experiences testing the system and demographic information.

Figure 6. Practice tactile overlays, mounted on a cardboard sheet, were shown to participants before testing each application.

FINDINGS

Task 1: Chart Comprehension

Chart Reading Time

We measured the time that each participant spent with each of the bar charts. The average time per chart by interface type is shown in Table 2.

Graph Interface Format	Mean Task Time (sec)	Mean Accuracy
Tactile	159.7 ($\sigma = 87.95$)	92.9%
Audio	214.0 ($\sigma = 80.84$)	92.9%

Table 2. Average time spent and accuracy with each dataset by condition.

A paired t-test was run on the task time observations to determine whether there was a statistically significant difference on the time spent on tactile charts vs. audio charts. There was no a significant difference between the two conditions ($t(13) = 1.769$, $p = .10$, Cohen's $d = 0.47$)

Chart Comprehension Questions

For each chart, participants answered five questions: identify the minimum, maximum, relative difference between points 1 and 2, relative difference between points 2 and 3, and

relative difference between points 3 and 4 (5 questions $\times$ 4 datasets = 20 responses per participant). The researchers coded each answer as correct or incorrect. The accuracy scores for each condition are shown in Table 2. Accuracy was the same for both conditions.

Subjective Preference

We asked participants whether they preferred interacting with bar charts via the tactile interface (i.e., RoboGraphics) or the audio-only interface. Participants rated their opinion on a Likert scale (1 = strongly prefer audio, 5 = strongly prefer tactile). Results are shown in Figure 7. All participants indicated that they preferred the tactile bar chart interface.

Figure 7. Number of participants by bar chart interface preference (1 = strongly prefer audio, 5 = strongly prefer tactile).

We also asked participants to rate each of the interfaces by perceived enjoyability and informativeness (Figure 8). On average, participants reported that the enjoyability of the tactile interface (4.43) was greater than that of the audio interface (3.14) and reported the tactile interface (4.43) to be more informative than the audio interface (3.57).

Figure 8. Average Likert responses and standard deviation for Enjoyability and Informativeness between the Graph Interface Conditions. (1 = least enjoyable/informative, 5 = most enjoyable/informative)

Task 2: Application Feedback

Following the second task, we asked participants to rate their favorite and least favorite of the proof-of-concept applications (Figure 9). The most favorite application was the tactile version of the Bar Charts application with 3 votes, followed by Tortoise and the Hare with 2 votes. The least favorite applications were the Analog Clock with 3 votes, followed by the audio-only Bar Charts application with 2 votes.

Figure 9. Reported favorite applications (left) and least favorite (right) applications presented in the study.

Participants also provided detailed comments about their experience using RoboGraphics. This feedback is summarized in the next section.

DISCUSSION

In this section we provide more context regarding *how* participants interacted with the various interfaces and summarize participants' feedback about their experiences.

Exploration Strategies

Participants applied different strategies when exploring the audio-only bar chart vs. the tactile bar chart.

With the audio bar charts, participants adopted two main strategies. Some participants (P1, P2, P3, P4, and P5) performed a bottom-up search, starting at the bottom of the screen and moving up until they found the data points. Other participants (P6 and P7) used side-to-side motions, locating one point and then using that point as a reference while searching for other points.

With the tactile bar charts, participants were able to perform some additional types of movements. In particular, most participants practiced touching the main part of the display with both hands in order to feel the overall trend of the chart. Sometimes participants used this two-handed grasp to gain a general sense of the trend but would then hold one robot as a reference point while touching the others.

Although we did not ask detailed questions about participants' familiarity with reading charts, some of our participants expressed uncertainty about their experience

with reading and interpreting graphical information. When discussing the bar charts task, P2 mentioned, "It's been so long since I looked at a graph, I forgot to think about a graph, and I thought it was just space." P4 noted that he had not explored a graph in several years. Conversely, P5, who had a background in economics, reported having extensive experience reading charts and graphs. It is likely that participants' experiences with graphical data may affect their ability to use RoboGraphics. We suspect that more practice with RoboGraphics may lead to improved performance but leave the answer to this question as an opportunity for future work.

Interactions with Moving Objects

One benefit of our mobile robot-based haptic display is the ability to represent and track moving objects [8]. This interaction was highlighted in applications such as Tortoise and the Hare, which centered around the relative movement of two objects.

We observed several different hand poses during use of the Tortoise and the Hare application. P1 and P2 used one hand to follow the Tortoise and a second hand to follow the Hare, while P3 tracked both characters with one hand. Both P1 and P3 frequently reached out to check on the tactile landmarks; P1 stretched his hand out to touch both the Hare and the landmarks, while P3 used his other hand to track the landmarks.

Both P3 and P4 noted that they considered the sounds made by the robots to provide an additional information channel. P4 described using the sounds made by the robots during the Tortoise and the Hare to track each of the characters. P3 described a similar strategy when exploring the tactile bar chart, moving the robots to the bottom of the screen and listening to how long it took for them to move into their correct positions.

While we lack sufficient data to identify which (if any) exploration strategies work best, we note that participants explored and identified multiple strategies for making sense of information displayed on the device.

Benefits of RoboGraphics

We compared the total amount of time participants spent exploring the bar charts across the two interface conditions, as well as their ability to correctly answer questions about the data. Overall, there was no significant performance difference between the two conditions. Participants were marginally (but non-significantly) faster using the tactile interface but answered the same number of comprehension questions correctly. Since we did not give participants a time limit, it is not surprising to see similar performance characteristics. Given the formative focus of this study, we are pleased to note that the additional complexity of the mobile robots did not adversely affect performance but did positively affect subjective preference.

Overall, participants preferred interacting with chart data using the combination of static tactile overlays and mobile

robots. When asked to describe why they preferred RoboGraphics, participants commented on both the speed of finding the robots, as well as the importance of *gestalt* understanding of the graphical information. P4 noted that, “With the robots, you don’t have to find the data points because they are already there.” Similarly, P1 said, “I found the tactile ones to be easier to locate”. P5 stated, “I like this better ... I can tell. You can visualize it and sense the distance and spatial orientation.”

Conversely, the audio-only interface did not seem to provide this awareness of the gestalt. P2 said “I think the robots ... that’s what we are missing here. [With the robots] I can put my hands like this and know where everything is.”

Participants noted additional benefits of RoboGraphics after testing the other proof-of-concept applications. Several participants noted that animating tactile graphics can sometimes result in simpler graphics, P1 stated, “Tactile graphics can get a little bit busy, this is nice because you can follow the track.” P3 recalled experiencing difficulty when first learning to read tactile graphics due to the need to frequently move between the graphic and the legend. P2 noted that, with RoboGraphics, “all the start and end points are all in the same spot [so] I have to scan less.”

Animating tactile graphics can also be useful in representing specific types of information. P5 mentioned that he sometimes uses tactile graphics to understand sequential processes but finds it difficult to move between the legend and multiple points on the diagram, whereas RoboGraphics could direct his hand to the correct location. P7, who teaches Braille, noted that RoboGraphics may be useful not only for showing Braille characters, but because showing the animated transitions between characters may teach students about the relationships between those characters.

While participants noted some advantages of dynamic tactile graphics, they also pointed out some cases in which traditional tactile graphics are superior. Specifically, RoboGraphics have a much lower resolution than traditionally produced tactile graphics. P1 noted that, “with the [static] tactile, you get a better view of the whole thing.” P7 reported that traditional tactile graphics are useful for precision tasks such as dense charts or detailed representations such as a clock face, which was somewhat difficult to read given the relatively large size of the robots.

Future Applications and Opportunities

We asked participants which of the applications they liked most, as well as what other applications they would be interested in seeing in the future.

Of the initial set of applications, the tactile bar chart was the most popular. Overall, participants seemed to agree that RoboGraphics could be a useful tool for exploring data. P5, who has a background in economics, felt that the display would have been useful for his graduate studies because it could convey the shape of data very quickly. When asked what he might use RoboGraphics for in the future, P5

requested support for more types of data visualizations, including pie charts, scatterplots, function graphs, and representations specific to economics such as supply-and-demand curves. P6 requested additional support for geographic data such as turn-by-turn map directions, topographical maps, and animations that demonstrate geological processes.

The animated story of the Tortoise and the Hare was also popular, and participants provided additional ideas for representing stories and games using RoboGraphics. P2 noted that she was smiling throughout the Tortoise and the Hare and said that she was eager to read more stories in this format. P2 thought that RoboGraphics could be used to present other fairy tales such as Hansel and Gretel, could represent riddles such as the river-crossing riddle in which a fox, a chicken, and a bag of grain need to be transported across a river, and could perhaps be extended to support narrative games such as The Oregon Trail.

Participants also suggested new applications for this technology. P3, who is a writer, suggested that RoboGraphics could be used to represent typefaces and could be used to teach topics in typography such as serif vs. sans-serif fonts. P3 also suggested that RoboGraphics could be used to teach him music notation symbols so that he could learn to write music. P4 noted that RoboGraphics could be used to create aesthetically pleasing displays, and that he might want to use the system as a decoration for his home.

LIMITATIONS AND FUTURE WORK

Although participants were generally positive about their experience using RoboGraphics, there are opportunities to improve the usability and accessibility of the system. The small robots used in the system present some usability challenges, such as turning slowly, accidentally being knocked off course, and shutting off to preserve power. At the same time, participants noted that the robots were sometimes too large to represent fine details. As robot-based shape displays are still relatively new, there are exciting opportunities in exploring the design space of the associated robot hardware and in developing robot platforms that are optimized for non-visual displays.

In developing RoboGraphics, we found that distinguishing robots from one another could sometimes be difficult. To address this problem, we developed tactile mounts that changed the shape of the robot. Moving from a swarm of indistinguishable robots to a set of robots with identifiable features may create new opportunities for tangible data displays. Future iterations of the tactile mount system could support robots of different sizes, shapes, and textures, or even robots with additional capabilities, such as the ability for multiple robots to link together.

A major limitation of the current RoboGraphics system is that content can only be authored by programmers. Future versions of RoboGraphics should support a variety of content authoring tools, supporting both programmers and less

technical subject experts such as tactile artists. Ideally these authoring tools should be usable by both blind and sighted users as well.

Finally, while the current version of RoboGraphics uses a small display and a small number of robots, it may be worth exploring how the number of robots affects the ability to represent certain types of information. Further empirical studies of how robot-based shape displays represent information could lead to guidelines for designing such displays and could also result in algorithms for mapping information onto shape displays of different sizes or with different numbers of robots.

CONCLUSION

In this paper, we introduced RoboGraphics, a system that combines a touch screen with static tactile overlays and mobile robots to create dynamic tactile graphics. Our prototype system supports a variety of applications including dynamic data graphics, audio-tactile stories, and interactive applications. Participants in our user study were able to read and understand data graphics created by our system, and they suggested a variety of future applications for this technology. While innovations in high-tech shape displays will likely lead to new advances in accessibility, our work demonstrates that there remain new ways to use existing technology to create more accessible media.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under grants IIS-1619384 and IIS-1652907. Any opinions, findings, conclusions or recommendations expressed in this work are those of the authors and do not necessarily reflect those of the National Science Foundation.

REFERENCES

- [1] Parastoo Abtahi and Sean Follmer. 2018. Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18). ACM, New York, NY, USA, Paper 150, 13 pages. DOI: <https://doi.org/10.1145/3173574.3173724>
- [2] Alphaboxes – Teaching braille characters and phonetic awareness of letter sounds using a six-compartment box, 2019. <https://www.tsbvi.edu/fall-2011-items/3308-alphaboxes>
- [3] 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, Article 3 (January 2016), 22 pages. DOI: <http://dx.doi.org/10.1145/2854005>
- [4] Emeline Brule, Gilles Bailly, Anke Brock, Frederic Valentin, Grégoire Denis, and Christophe Jouffrais. 2016. MapSense: Multi-Sensory Interactive Maps for Children Living with Visual Impairments. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16). ACM, New York, NY, USA, 445-457. DOI: <https://doi.org/10.1145/2858036.2858375>
- [5] Julie Ducasse, Marc J-M Macé, Marcos Serrano, and Christophe Jouffrais. 2016. Tangible Reels: Construction and Exploration of Tangible Maps by Visually Impaired Users. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16). ACM, New York, NY, USA, 2186-2197. DOI: <https://doi.org/10.1145/2858036.2858058>
- [6] Aluna Everitt and Jason Alexander. 2017. PolySurface: A Design Approach for Rapid Prototyping of Shape-Changing Displays Using Semi-Solid Surfaces. In *Proceedings of the 2017 Conference on Designing Interactive Systems* (DIS '17). ACM, New York, NY, USA, 1283-1294. DOI: <https://doi.org/10.1145/3064663.3064677>
- [7] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. 2013. inFORM: dynamic physical affordances and constraints through shape and object actuation. In *Proceedings of the 26th annual ACM symposium on User interface software and technology* (UIST '13). ACM, New York, NY, USA, 417-426. DOI: <https://doi.org/10.1145/2501988.2502032>
- [8] Darren Guinness, Annika Muehlbradt, Daniel Szafr, and Shaun K. Kane. 2018. The Haptic Video Player: Using Mobile Robots to Create Tangible Video Annotations. In *Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces* (ISS '18). ACM, New York, NY, USA, 203-211. DOI: <https://doi.org/10.1145/3279778.3279805>
- [9] Darren Guinness, Daniel Szafr, and Shaun K. Kane. 2017. GUI Robots: Using Off-the-Shelf Robots as Tangible Input and Output Devices for Unmodified GUI Applications. In *Proceedings of the 2017 Conference on Designing Interactive Systems* (DIS '17). ACM, New York, NY, USA, 767-778. DOI: <https://doi.org/10.1145/3064663.3064706>
- [10] How tech helped blind students 'watch' the solar eclipse, 2017. <https://www.cnet.com/news/how-tech-helped-blind-students-watch-the-solar-eclipse/>
- [11] Joseph Jacobs, ed. *The fables of Aesop*. Hurst, 1902.
- [12] 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, USA, 273-282. DOI: <https://doi.org/10.1145/2047196.2047232>
- [13] Shaun K. Kane, Meredith Ringel Morris, and Jacob O. Wobbrock. 2013. Touchplates: low-cost tactile overlays for visually impaired touch screen users. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '13). ACM,

- New York, NY, USA, , Article 22 , 8 pages. DOI: <http://dx.doi.org/10.1145/2513383.2513442>
- [14] Lawrence H. Kim and Sean Follmer. 2019. SwarmHaptics: Haptic Display with Swarm Robots. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19). ACM, New York, NY, USA, Paper 688, 13 pages. DOI: <https://doi.org/10.1145/3290605.3300918>
- [15] Lawrence H. Kim and Sean Follmer. 2017. UbiSwarm: Ubiquitous Robotic Interfaces and Investigation of Abstract Motion as a Display. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 1, 3, Article 66 (September 2017), 20 pages. DOI: <https://doi.org/10.1145/3130931>
- [16] Nicolai Marquardt, Miguel A. Nacenta, James E. Young, Sheelagh Carpendale, Saul Greenberg, and Ehud Sharlin. 2009. The Haptic Tabletop Puck: tactile feedback for interactive tabletops. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '09)*. ACM, New York, NY, USA, 85-92. DOI: <https://doi.org/10.1145/1731903.1731922>
- [17] Ayberk Özgür, Séverin Lemaignan, Wafa Johal, Maria Beltran, Manon Briod, Léa Pereyre, Francesco Mondada, and Pierre Dillenbourg. 2017. Cellulo: Versatile Handheld Robots for Education. In Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI '17). ACM, New York, NY, USA, 119-127. DOI: <https://doi.org/10.1145/2909824.3020247>
- [18] Mathieu Le Goc, Lawrence H. Kim, Ali Parsaei, Jean-Daniel Fekete, Pierre Dragicevic, and Sean Follmer. 2016. Zooids: Building Blocks for Swarm User Interfaces. In Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16). ACM, New York, NY, USA, 97-109. DOI: <https://doi.org/10.1145/2984511.2984547>
- [19] Lei Shi, Yuhang Zhao, and Shiri Azenkot. 2017. Markit and Talkit: A Low-Barrier Toolkit to Augment 3D Printed Models with Audio Annotations. In Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17). ACM, New York, NY, USA, 493-506. DOI: <https://doi.org/10.1145/3126594.3126650>
- [20] Abigale Stangl, Jeeun Kim, and Tom Yeh. 2014. 3D printed tactile picture books for children with visual impairments: a design probe. In Proceedings of the 2014 conference on Interaction design and children (IDC '14). ACM, New York, NY, USA, 321-324. DOI: <https://doi.org/10.1145/2593968.2610482>
- [21] Ryo Suzuki, Abigale Stangl, Mark D. Gross, and Tom Yeh. 2017. FluxMarker: Enhancing Tactile Graphics with Dynamic Tactile Markers. In Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '17). ACM, New York, NY, USA, 190-199. DOI: <https://doi.org/10.1145/3132525.3132548>
- [22] The Magic School Bus – Scholastic, 2019. <https://kids.scholastic.com/kids/books/the-magic-school-bus/>
- [23] TinyBop Mammals – Mammals education app for kids by TinyBop, 2019. <https://tinybop.com/apps/mammals>
- [24] Brianna J. Tomlinson, R. Michael Winters, Christopher Latina, Smruthi Bhat, Milap Rane, & Bruce N. Walker. 2017. Solar System Sonification: Exploring Earth and its Neighbors through Sound Georgia Tech Center for Music Technology (GTCMT) 2 , Georgia Institute of Technology , Atlanta , USA, (Icad), 128–134.
- [25] Yu, Wai, and Stephen Brewster. "Comparing two haptic interfaces for multimodal graph rendering." In Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2002, pp. 3-9. IEEE, 2002.
- [26] Steven Wall and Stephen Brewster. 2006. Feeling what you hear: tactile feedback for navigation of audio graphs. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06), Rebecca Grinter, Thomas Rodden, Paul Aoki, Ed Cutrell, Robin Jeffries, and Gary Olson (Eds.). ACM, New York, NY, USA, 1123-1132. DOI=<http://dx.doi.org/10.1145/1124772.1124941>
- [27] YouDescribe – Audio Descriptions for YouTube videos, 2019. <https://youdescribe.org>