

Usable Gestures for Blind People: Understanding Preference and Performance

Shaun K. Kane, Jacob O. Wobbrock

The Information School, DUB Group
University of Washington
Seattle, WA 98195 USA
{skane, wobbrock}@uw.edu

Richard E. Ladner

Computer Science & Engineering, DUB Group
University of Washington
Seattle, WA 98195 USA
ladner@cs.washington.edu

ABSTRACT

Despite growing awareness of the accessibility issues surrounding touch screen use by blind people, designers still face challenges when creating accessible touch screen interfaces. One major stumbling block is a lack of understanding about how blind people actually use touch screens. We conducted two user studies that compared how blind people and sighted people use touch screen gestures. First, we conducted a gesture elicitation study in which 10 blind and 10 sighted people invented gestures to perform common computing tasks on a tablet PC. We found that blind people have different gesture preferences than sighted people, including preferences for edge-based gestures and gestures that involve tapping virtual keys on a keyboard. Second, we conducted a performance study in which the same participants performed a set of reference gestures. We found significant differences in the speed, size, and shape of gestures performed by blind people versus those performed by sighted people. Our results suggest new design guidelines for accessible touch screen interfaces.

Author Keywords: Accessibility, blind, touch screens, gestures, gesture recognition.

ACM Classification Keywords: H.5.2. Information interfaces and presentation: User Interfaces—*input devices and strategies*. K.4.2. Computers and Society: Social issues—*assistive technologies for persons with disabilities*.

General Terms: Experimentation, Human Factors.

INTRODUCTION

While touch screens were once rare, touch screen-based interfaces are now present across a wide range of everyday technologies, including mobile devices, personal computers, and public kiosks. As touch screens have become mainstream, it is crucial that touch screen-based interfaces be usable by people with all abilities, including blind and visually impaired people.

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

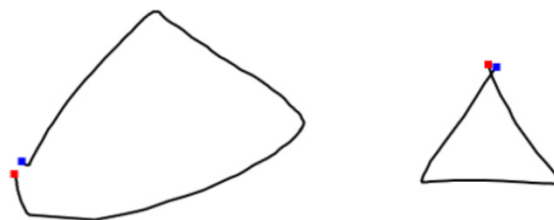


Figure 1. Two representative versions of a *triangle* gesture produced by a blind person (left) and a sighted person (right).

Until recently, most touch screens provided few or no accessibility features, leaving them largely unusable by blind people. However, both the blind community and technology manufacturers have made progress on this issue in recent years. At the 2009 Consumer Electronics Show, musician Stevie Wonder and a group of blind engineers took designers to task for the inaccessibility of touch screens, and encouraged them to improve touch screen accessibility for blind and visually impaired people [7]. Later that year, both Google and Apple released basic screen-reading software for their touch screen-based mobile devices, and most Google Android and Apple iOS devices now ship with screen-reading software preinstalled.

However, accessible touch screens still present challenges to both users and designers. Users must be able to learn new touch screen applications quickly and effectively, while designers must be able to implement accessible touch screen interaction techniques for a diverse range of devices and applications. Because most user interface designers are sighted, they may have a limited understanding of how blind people experience technology. We therefore argue that accessible touch screen interfaces can be improved substantially if *designers* can better understand how blind people actually use touch screens (Figure 1).

A designer who wishes to create a new accessible touch screen-based application currently faces several challenges. First, while touch screen interfaces for sighted users are largely consistent due to now-familiar gestures such as tapping, swiping, and pinching, touch screen interfaces for blind users vary widely across platforms. For example, in Apple's VoiceOver for iPhone¹, a user can browse a menu by performing a series of discrete directional flicks to move the cursor, and can select a menu item by double-tapping

¹ <http://www.apple.com/accessibility/iphone/vision.html>

the screen. In contrast, in Google's Eyes-Free Shell for Android², a user can select a menu item by holding their finger down on the screen, dragging their finger in one of 8 directions to choose the item, and then releasing their finger to select it. Both systems use completely different layouts and interaction primitives, and thus there is currently no *lingua franca* for touch screen interactions for blind people.

Second, there exist very few examples of how to extend accessible touch screen interfaces to devices other than smartphones. Touch screens may appear on devices of many different sizes, from jewelry-sized displays (e.g., [3]) to wall-sized interactive installations (e.g., [12]). However, most commercially available and accessible touch screen-based devices are smartphones. There still exists little information about how to best design accessible touch screen interfaces for tablets and other large touch screens.

Third, a designer who wishes to provide gestures in their application must consider whether the gestures will be appropriate for a blind user. Although blind people may use the same hardware as their sighted peers, it is possible that they will prefer to use different gestures, or that they will perform the same gestures differently than a sighted person. Sighted people perform gestures differently when they lack visual feedback [25], and it is reasonable to assume that a blind person may also perform gestures differently than a sighted person.

These challenges raise fundamental questions about how blind people use touch screens: What types of gestures are the most intuitive and easy to perform for a blind person? And if blind people perform gestures differently than sighted people, how do their gestures differ? In this paper, we address these questions through two user studies that explore how blind and sighted people interact with touch screens. First, we conducted a gesture elicitation study (see [28]) in which blind and sighted participants invented gestures for performing common computing tasks on a touch screen-based tablet PC. Second, we conducted a gesture performance study in which both blind and sighted participants repeatedly performed a set of standard gestures on a touch screen. Our results show that there are indeed differences in the types of gestures preferred by blind and sighted people, as well as differences in how gestures are performed by blind and sighted people. Based on these results, we provide suggestions for the design of future touch screen-based applications and devices for blind users.

RELATED WORK

Our research is motivated by prior attempts to create accessible touch screen user interfaces for blind people, as well as by prior studies of spatial and tactile perception in blind people. We also adopt participatory techniques from prior research to elicit participants' preferred gestures.

Gesture-Based User Interfaces for Blind People

Providing blind people with access to touch screens has been a concern since the creation of the earliest touch screen systems. Buxton *et al.*'s introduction of an early touch tablet [6] was followed shortly by a panel discussing accessibility issues surrounding touch interfaces [5]. In the 1990s, the emergence of touch screen kiosks in public places such as airports and shopping malls prompted investigation of how touch screen hardware could be made more accessible (e.g., [17,27]). In recent years, researchers have explored accessible interaction techniques for mobile touch screens (e.g., [4,11,15,22,23]), and commercial manufacturers have begun to incorporate screen-reading software into their mobile devices (e.g., VoiceOver¹, Eyes-Free Shell², and Mobile Speak for Windows Mobile³).

While some earlier systems such as the Talking Fingertip Technique [27] and the Talking Tactile Tablet [16] used physically adapted hardware, most of these systems have used software alone to enable accessible interactions on a touch screen, typically by accepting gestures as input and providing speech and audio as output. Despite the diversity of the touch screen-based devices that have been adapted for use by blind people, most of these systems have used one of a small set of underlying interaction techniques. As no formal taxonomy of these interfaces has yet been published, we refer to these techniques as *menu browsing*, *discrete gestures*, and *fixed regions*. These techniques are described briefly below.

Menu Browsing

In *menu browsing*, the user moves a cursor through a list of menu items and receives speech or audio feedback describing each item. The user then performs a gesture to actuate the currently selected item. The user may move the cursor through continuous touch gestures, in which the user strokes their finger across the screen to navigate the list, or by using discrete gestures or taps to move the cursor. The list of menu items typically changes based on the current application state. Menu items may be laid out spatially across the surface of the screen, either in their original arrangement or rearranged to optimize non-visual exploration. Systems that use *menu browsing* include the Talking Fingertip Technique [27], Slide Rule [15], EarPod [31], NavTouch [11], No-Look Notes [4], and VoiceOver¹.

Discrete Gestures

Other applications use a series of predefined *discrete gestures* to perform actions. Common actions are associated with a specific, predefined gesture, such as swiping one's finger in a specific direction or drawing a shape gesture. For example, swiping a finger from the top of the bottom of the screen may change the currently playing track in a touch screen-based music player. Systems that use *discrete gestures* include the Adaptive Blind Interaction Technique for Touchscreens [29], mBN [23], McGookin *et al.*'s gestural music player [19], and Eyes-Free Shell².

² <http://code.google.com/p/eyes-free>

³ <http://codefactory.es>

Fixed Regions

Some accessible touch screen applications map specific regions of the screen to predefined functions, as if the screen were in fact a set of discrete hardware buttons. In these applications, the user initiates an action by performing a gesture over the designated screen region. For example, a user might double tap the lower left corner of the screen to start an application. Systems that use *fixed regions* include Mobile Messenger for the Blind [22] and Mobile Speak³.

Our research takes a technique-agnostic approach to understanding touch screen interactions for blind people. In our gesture elicitation study, participants could choose any type of gesture they wished to perform a given command. The results of this research show support for using a variety of touch screen interaction techniques.

Spatial Perception and Drawing Among Blind People

As touch screen interfaces become more common, it is important to provide equal access to these interfaces for blind people. However, in some cases, touch screen-based interfaces may in fact be preferable to interfaces that use fixed buttons, even for blind people. Prior research shows that blind people, even those who are born blind, may have substantial spatial and tactile abilities. Blind people use the regions of the brain designated for visual processing when reading Braille and performing other spatial tasks [21]. Other studies have shown that both *early-blind* (those blind at birth or at an early age) and *late-blind* people have higher tactile sensitivity in their fingers than sighted people [10,26], and that late-blind adults can trace tactile shapes faster and more accurately than sighted adults [13].

Several studies have also examined the ability of blind people to draw or trace shapes using various technologies. Kamel and Landay [14] performed a study in which blind users drew shapes using both tactile swell paper and a keyboard-driven, grid-based drawing program. Crossan and Brewster [8] combined a haptic controller with audio feedback to enable blind users to trace simple shapes.

This research confirms that blind people are capable of performing gestures and drawing shapes on a screen, and suggests that gestures may in fact be an effective interaction method for some blind people.

Participatory Design of Gestures

Another limitation of current accessible touch screens is the lack of user participation in their development. Although some researchers have incorporated user feedback into the design of accessible touch screen systems (*e.g.*, [4,15]), the resulting systems have still been largely created by their designers. However, past research has shown that users often prefer gestures that were created by groups of potential users to those created by a single designer [20]. Beaudouin-Lafon [2] and Liu *et al.* [18] each created gesture sets based on users' observed movements, and Epps *et al.* [9] elicited gestures from users by showing them images of computing tasks. Wobbrock *et al.* [28] used a

participatory design approach to create a user-defined gesture set. In that study, participants were shown the outcome of an action and asked to demonstrate gestures that would accomplish that action.

The gesture elicitation study presented here is based on methods introduced by Wobbrock *et al.* [28]. However, while that study focused on visual touch screen interactions for sighted users, the current study examines gesture preferences among both blind and sighted participants.

STUDY 1: GESTURE ELICITATION

To better understand how blind people might prefer to interact with a touch screen, we conducted a study in which both blind and sighted participants were asked to invent gestures that could be used to conduct standard computing tasks on a touch screen-based tablet PC.

Participants

We recruited 10 blind people (6 female, 4 male, average age 49.0, SD=12.2) and 10 sighted people (4 female, 6 male, average age 28.8, SD=11.6) for the study. For the purposes of this research, we define blind participants as people who typically use a screen reader to access a computer. Blind participants were recruited via local blind organizations and via word of mouth. Sighted participants were recruited via local mailing lists and bulletin boards. Nine of the sighted participants and 6 of the blind participants regularly used some touch screen-based device. All of the sighted participants and 8 of the blind participants were right-handed. Three of the blind participants were early-blind and had become blind by the age of 2 years old.

Apparatus

Participants executed the experimental tasks using a 10.1-inch Lenovo S10-3t multi-touch tablet PC (Figure 2). The tablet PC ran Windows 7 and a custom C# application that recorded all screen touches with millisecond timestamps. Touch information was stored in a JSON-formatted log file. The experimenter used a wireless keyboard to begin and end experimental trials. A video camera captured the tablet PC and the participants' hands, as well as participants' spoken comments and think-aloud data.

Procedure

The study protocol was based upon the user-defined gesture study by Wobbrock *et al.* [28], but was modified to be used with both blind and sighted participants. Participants were seated at a desk in front of the tablet PC. The tablet PC was initialized to a blank screen, and participants were introduced to the device. Blind participants were given an opportunity to touch the screen and bezel in order to orient themselves. Sighted participants were shown that the screen tracked touches and visualized them as trails on the screen.

Once participants were ready to proceed, the experimenter began the session. At this point the application's blank screen was replaced by a "shapes world" containing labeled squares and circles. This neutral layout was chosen, as in Wobbrock *et al.* [28], to remind participants that they were

not creating gestures for any specific operating system or application. The shapes on the screen provided both visual and audio feedback: for example, touching a square would cause the program to speak the word “square,” and a white noise loop played while the participant held his or her finger over the shape.



Figure 2. A participant performs an experimental task using the tablet PC.

The experimenter informed participants that their task was to invent gestures that could be used to execute a set of computing commands. Commands were derived from Wobbrock *et al.* [28] and Morris *et al.* [20]. However, commands that had a primarily visual function (*e.g.*, enlarge, zoom in) were replaced by commands that applied to both visual and non-visual interfaces (*e.g.*, move up in hierarchy). The commands used in the study were: *menu*, *help*, *undo*, *task switch*, *move down in hierarchy*, *move up in hierarchy*, *previous page*, *next page*, *accept*, *reject*, *move object*, *open*, *close*, *duplicate*, *delete*, *cut*, *paste*, *select single*, *select group*, *move insertion point*, *select text range*, and *enter text*.

Because showing the outcome of the command visually would not be accessible to all participants, the experimenter read a standardized description that described the *outcome* of the command. For example, for the *next page* command, the participant was told: “Next page. You move from the current page of content to the next page.” Commands were presented to each participant in random order.

For each trial, the experimenter read the description for the command. The participant was then asked to invent 2 different gestures that could initiate the command, and to think aloud while doing so. Once the participant decided upon a gesture, they described the gesture verbally to the experimenter and demonstrated it 3 times using the tablet PC’s touch screen. Once the participant had demonstrated the 2 gestures that they had invented, the experimenter prompted them to rate each of the gestures using scales from Wobbrock *et al.* [28]. The first scale, referred to here as *good match*, read: “The gesture I picked is a good match for its intended purpose.” The second scale, referred to here as *easiness*, read: “The gesture I picked is easy to perform.” Both questions used Likert-type scales that ranged from 1 (strongly disagree) to 7 (strongly agree).

Results

Each participant invented 2 gestures for each of the 22 commands. We collected a total of $20 \times 22 \times 2 = 880$ gestures. In this section we analyze differences between the gestures invented by blind participants and the gestures invented by sighted participants, including participants’ preferred methods for entering text on a touch screen.

Gesture Ratings

For the *good match* question, blind participants gave the gestures they created an average score of 5.54 (SD=1.15), and sighted participants gave the gestures they created an average score of 5.15 (SD=1.42). For the *easiness* question, blind participants gave their gestures an average score of 5.77 (SD=1.15), while sighted participants gave their gestures an average score of 5.72 (SD=1.39). Logistic regression showed that blind participants rated the gestures they created as significantly better on the *good match* question ($\chi^2_{(1,N=880)}=13.69$, $p<.001$). There was no significant difference between the two groups on the *easiness* question ($\chi^2_{(1,N=880)}=0.23$, *n.s.*).

Gesture Properties

We examined differences in the properties of the gestures invented by blind and sighted participants, including the total number of strokes, the location of the gesture, and the use of multi-touch.

Stroke count. On average, blind participants’ gestures used 1.52 (SD=0.75) strokes and sighted participants’ gestures used 1.24 (SD=0.46) strokes. A Wilcoxon rank-sum test revealed that blind participants’ gestures contained significantly more strokes ($Z=5.63$, $N=880$, $p<.0001$).

Location. We analyzed whether participants’ gestures used either an edge or corner of the screen. Of the 880 gestures produced, 113 (12.8%) used an edge and 213 (24.2%) used a corner. Of these, blind people invented 69 (61.1%) of the edge gestures and 138 (64.8%) of the corner gestures. A Chi-Square test showed that blind people were more likely to choose gestures that used the edge ($\chi^2_{(1,N=880)}=6.35$, $p<.05$) or corner ($\chi^2_{(1,N=880)}=24.58$, $p<.0001$).

Multi-touch. Of the 880 gestures, 270 (30.7%) used at least two simultaneous contact points at some point during the gesture. Blind participants invented 166 (61.5%) of these multi-touch gestures. A Chi-Square test showed that blind participants were significantly more likely than sighted participants to invent multi-touch gestures ($\chi^2_{(1,N=880)}=20.54$, $p<.0001$). Although both blind and sighted participants invented multi-touch gestures, the groups used multi-touch differently, and many of the multi-touch gestures invented by blind participants were different than the multi-touch gestures used in current gesture-based user interfaces. In particular, many of the multi-touch gestures performed by blind participants involved a virtual mode key, in which the participant held one finger down on a specific area of the screen while performing the gesture with a second finger or hand. While thinking aloud, some blind participants

justified the first touch as a way to reduce potential conflicts between gestures. This use of a mode key was common among blind participants: 63 of the gestures invented by blind participants during the study used a mode key to activate the gesture, while only 10 of the gestures invented by sighted people used a mode key.

Gesture Nature and Rationale

We identified participants' rationale for the gestures that they created by analyzing their think-aloud comments and gesture descriptions. We classified gesture rationale using the *nature* dimension from Wobbrock *et al.*'s taxonomy of gestures [28]. Gesture nature considers a gesture's underlying explanation as either *symbolic* ("gesture visually depicts a symbol"), *physical* ("gesture acts physically on objects"), *metaphorical* ("gesture indicates a metaphor"), or *abstract* ("gesture-referent mapping is arbitrary"). We assigned a nature to each gesture produced during the gesture elicitation study. Figure 3 shows the gesture natures produced during this study.

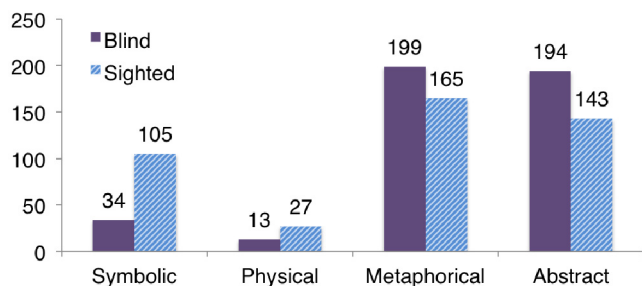


Figure 3. Gesture rationale for each subject group as described using gesture nature.

A Chi-Square test showed a significant difference between gesture natures assigned by blind and sighted participants ($\chi^2_{(3, N=880)}=52.06, p<.0001$). Examining the results more closely, we note a higher number of symbolic gestures invented by sighted participants, and a higher number of abstract and metaphorical gestures invented by blind participants. Within these categories, participants in both groups provided similar explanations for their gestures, with one notable exception: 95 of the metaphorical gestures produced by blind participants involved an interaction in which the participant touched areas of the screen in a way that was analogous to pressing keys on a physical keyboard. For example, one participant demonstrated the paste command with a CONTROL-V gesture, in which she tapped an area of the screen near where the CONTROL key would be on a QWERTY keyboard, and then tapped an area of the screen near where the V key would be. These keyboard metaphors accounted for 21.6% of the gestures invented by blind participants during the study, and were used at least once by 9 of the 10 blind participants. Of the 9 blind participants who used such a gesture, 5 regularly used some touch screen-based device, suggesting that this idea was popular even among participants who were familiar with traditional touch screen interfaces. No sighted participants performed gestures based on a physical keyboard layout.

Preferences for Text Entry

We included the *enter text* command in this study to elicit ideas about how a user might enter text using a touch screen. We were particularly interested in suggestions for text entry methods from blind participants, as entering text without visual feedback can be slow and laborious. Participants offered the following ideas for entering text:

- *On-screen QWERTY keyboard*: suggested by 7 blind and 9 sighted participants;
- *Handwriting*: Suggested by 2 blind and 3 sighted participants;
- *Perkins Braille*: Two blind participants suggested using the multi-touch screen to enter Braille using the Perkins technique, a two-handed chording technique used on physical Braille typewriters;
- *T9³*: Two blind participants suggested using the T9 predictive text method found on 12-key phone keypads.

Methods suggested by only one participant were: tapping out the dots of Braille characters, an on-screen Dvorak keyboard, Graffiti, and ShapeWriter (*a.k.a.* SHARK) [30]. Although participants were asked to keep an open mind and be creative, 10 participants were unable to invent a second text entry method, and instead chose an alternative form of their original method. Of these participants, 9 chose a variation of QWERTY and 1 chose a variation of T9.

STUDY 2: GESTURE PERFORMANCE

The gesture elicitation study provided us with new insight about the types of gestures that blind people may wish to perform. However, the open-ended nature of the study made it difficult to ascertain whether blind people and sighted people actually perform gestures differently, or whether they simply prefer different types of gestures. To determine if there were significant differences in how blind and sighted people performed the *same* gestures, we conducted a second study in which all participants performed the same set of standard gestures.

Participants and Apparatus

This study featured the same participants as the previous gesture elicitation study. Participants used the same Lenovo tablet PC and logging software as the previous study. The tablet PC ran the same logging application, which recorded all contacts with millisecond timestamps. During each trial, the application initially showed a blank screen, and drew trails that visualized the user's touches on the screen.

Procedure

The experimental procedure was similar to the gesture elicitation study. However, instead of inventing new gestures, participants performed specific gestures as specified by the experimenter. For each gesture, the experimenter read the name and a brief description of the gesture. The participant was given a chance to practice the

⁴ <http://www.tegic.com>

gesture once on the touch screen. Once the participant had practiced the gesture, they performed the gesture 3 times, and each trial was recorded in the log file. After completing each gesture, the participant rated the gesture using a variation of the *easiness* scale from the first study: “The gesture is easy to perform.” (1 = strongly disagree, 7 = strongly agree).

For this study, we chose 40 gestures that represented common interactions on current touch screen platforms, and which included unistroke, multi-stroke, and multi-touch gestures. Gestures were divided into 5 categories: *tap* (taps in various locations on screen), *flick* (directional swiping gestures), *multi-touch* gestures, *shape* (simple geometric shapes), and *symbol* (letters, numbers, and other symbols). The following categories and gestures were used:

- *Tap*: single tap center, single tap left, single tap right, single tap top, single tap bottom, single tap top left, single tap top right, single tap bottom left, single tap bottom right, double tap, triple tap;
- *Flick*: flick left, flick right, flick up, flick down;
- *Multi-touch*: 2-finger pinch, 2-finger spread apart, 2-finger rotate clockwise, 2-finger rotate counterclockwise;
- *Shape*: square, circle, triangle;
- *Symbol*: A, B, C, D, E, F, L, X, Z, 1, 2, 3, 4, 5, question mark, check mark, 5-pointed star, scratch out.

The gesture set contained some symbols used in printed English writing, including the numbers 1 through 5 and the letters A through F, L, X, and Z. Although many blind people do not typically use handwriting, and thus may be unfamiliar with these symbols, we chose to include the symbols to increase the overall number of glyph-like gestures, as well as to explore how familiar blind people actually are with these symbols. Participants were allowed to skip a gesture if they were not familiar with it.

Results

Gestures Collected

Each participant was asked to perform 40 gestures 3 times each, for a total of $40 \times 3 \times 20 = 2400$ gestures performed. However, 2 blind participants skipped a total of 15 gestures because they were unfamiliar with the gesture. One participant skipped *A*, *E*, *2*, *4*, *5*, *check mark*, and *5-pointed star*, while the other skipped *B*, *Z*, *2*, *3*, *4*, *5*, *question mark*, and *5-pointed star*. In addition, 20 gestures were not accurately captured by the touch screen, leaving a total of 2335 recorded gesture instances and 785 gesture ratings.

Gesture Ratings

Participants rated each of the gestures in terms of *easiness*. Overall, blind participants gave the gestures they performed an average score of 5.71 (SD=1.55), while sighted participants gave the gestures they performed an average score of 5.76 (SD=1.28). Ratings by gesture category are shown in Figure 4. Logistic regression showed no

significant effects of blindness ($\chi^2_{(1,N=785)}=2.70$, *n.s.*) on *easiness*. However, there was a significant effect of gesture category ($\chi^2_{(4,N=785)}=68.85$, $p<.0001$) on *easiness*, which shows that participants’ ratings were influenced by the gesture category. There was also a significant interaction between blindness and gesture category ($\chi^2_{(4,N=785)}=18.61$, $p<.001$), indicating that a gesture’s category affected its rating differently for blind and sighted participants.

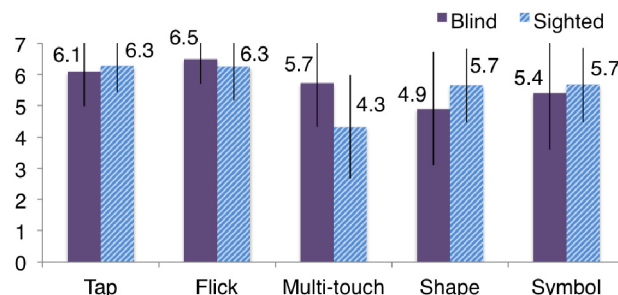


Figure 4. Participants’ *easiness* ratings by gesture category. Error bars indicate ± 1 SD.

Gesture Properties

In addition to soliciting ratings for each gesture, we examined various properties of the gestures that participants performed to determine whether there were significant differences in how blind and sighted people performed the same gestures.

Size. We measured the overall size of each gesture using the area of the bounding box for that gesture. The average size of gestures produced by blind participants was 165.82 (SD=136.89) pixels², while the average size of gestures produced by sighted participants was 148.05 (SD=120.68) pixels². A Wilcoxon rank-sum test found a significant difference in the size of gestures produced by blind and sighted participants ($Z=3.20$, $N=2335$, $p<.01$), indicating that blind participants tended to create significantly larger gestures than sighted participants.

Size variation. In addition to calculating gesture size, we also examined the size variation between multiple instances of the same gesture created by a single participant. We calculated size variation using the standard deviation of the gesture size for a given participant and gesture. This standard deviation averaged 20.43 (SD=23.38) pixels² for blind participants and 13.59 (SD=18.44) pixels² for sighted participants. There was a significant difference in the standard deviation of sizes between blind and sighted participants ($Z=5.32$, $N=780$, $p<.001$), indicating that blind participants produced gestures with greater variation in size when performing the same gesture multiple times.

Aspect ratio. We calculated the aspect ratio of each gesture’s bounding box (width/height). The average aspect ratio for blind participants was 1.64 (SD=3.73), and the average aspect ratio for sighted participants was 1.44 (SD=3.45), suggesting that blind participants tended to create wider gestures. However, when examining all

gestures performed in the study, we found no significant difference in the aspect ratio of gestures produced by blind and sighted participants ($Z=1.75$, $N=2335$, $p=.08$, *n.s.*). Looking specifically at gestures in which participants drew some glyph (*i.e.*, *shape* and *symbol* gestures), we did find that blind participants drew significantly wider gestures than sighted participants ($Z=3.87$, $N=1213$, $p<.0001$).

Speed. On average, blind participants performed gestures in 925.90 (SD=1110.22) milliseconds, while sighted participants performed gestures in 462.58 (SD=477.24) milliseconds. This difference was significant ($Z=10.96$, $N=2335$, $p<.0001$), showing that, on average, blind participants took approximately twice as long to perform the same gestures.

Analysis of Specific Gesture Features

In addition to the aforementioned properties, we were interested in some specific issues that had been observed during our pilot testing. Because these issues did not apply across all gestures, we examined these features using specific subsets of the gesture set, rather than the entire set.

Location accuracy. In general, participants were able to perform gestures at any location on the screen. However, nine gestures in the *tap* category referred to specific screen locations, including the corners, edges, and center. For these gestures, we calculated the distance between the centroid of the performed gesture and the specified screen location. For blind participants, the average distance from the specified location was 110.97 (SD=50.28) pixels. The average distance for sighted participants was 48.84 (SD=18.35) pixels. Blind participants had a significantly greater offset than sighted participants ($Z=-15.52$, $N=528$, $p<.0001$). During the study, some blind participants mentioned that it was difficult to target locations that were away from the screen corners. Looking only at gestures in which participants tapped the corners, we found that blind people's gestures were still farther from the intended targets than sighted people's gestures ($Z=-9.82$, $N=233$, $p<.0001$).

Form closure. Another issue that we observed during initial testing was the difficulty of connecting the various parts of a gesture without visual feedback. However, it is difficult to operationalize this feature for an arbitrary gesture, as different gestures connect in different ways.

To examine this issue quantitatively, we measured its effects on the *circle* gesture alone. We chose the circle gesture because it was likely to be completed in a single stroke, and because the start point and end points coincide in its canonical form. For the purposes of this analysis, we defined the metric *form closure distance* as the Cartesian distance between the circle's start point and end point. For blind participants, the average form closure distance was 149.12 (SD=84.00) pixels, while for sighted participants this distance was 67.12 (SD=94.59) pixels. A Wilcoxon rank-sum test revealed that this difference was statistically significant ($Z=4.47$, $N=60$, $p<.0001$). This suggests that

gestures created by blind people were more likely to have start and end points that did not coincide (Figure 5).

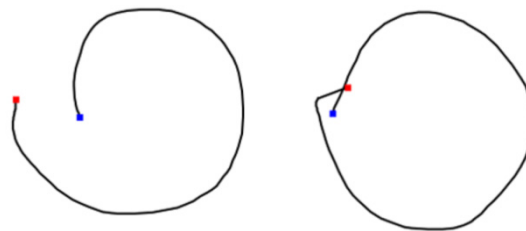


Figure 5. Two circle gestures drawn by a blind participant with high form closure distance, *i.e.*, a large distance between the start and end points.

Line steadiness. During pilot testing, we noted that the lines of some blind participants' gestures seemed to be less steady or more "wavy" than those created by sighted participants. Prior studies have attempted to quantify the steadiness of a gesture by measuring parameters of a single specific gesture, such as angular deviation when drawing a straight line or eccentricity from a reference shape (*e.g.*, [24]). However, these approaches are limited to specific gestures and do not generalize. Furthermore, because participants in this study could not see a reference shape on screen, matching their gestures to a reference shape would be inappropriate.

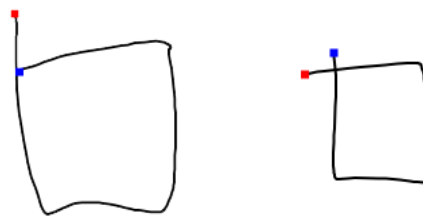


Figure 6. Lines drawn by blind participants tended to be less steady (left) than lines drawn by sighted participants (right).

To model this waviness quantitatively, we introduce a generalizable stability metric for drawn gestures, the *average angular acceleration* metric, shown in Equation 1.

$$A = \frac{\sum_p \frac{d\omega}{dt}}{|p|} \quad (1)$$

The average angular acceleration metric is defined as the sum of the angular acceleration over the course of the gesture, divided by the length of the gesture in points. This metric approximately measures how much the path changes direction over the course of the entire gesture. For example, a gesture that is drawn in a wavy or jagged fashion will continually be changing direction, and thus will have a higher average angular acceleration value. Figure 6 shows a wavy gesture with high average angular acceleration and a steady gesture with low average angular acceleration.

Note that this metric is not appropriate for comparing between different gestures, as each gesture requires a different amount of angular movement to be drawn

correctly. However, examining different participants as they perform the same gesture allows us to approximate the overall steadiness or waviness of their lines using this measure. For demonstration, we calculated this value for the *square* gesture only. For blind participants, the average angular acceleration was 0.47 (SD=0.26), while for sighted participants the average angular acceleration was 0.06 (SD=0.13). This difference was statistically significant ($Z=5.27$, $N=60$, $p<.0001$).

Gesture Recognition Accuracy

We have described a number of differences in how blind people and sighted people perform gestures. It seems likely that these differences would also affect gesture recognition accuracy for gestures performed by a blind person. To explore this question, we compared gesture recognition results for a subset of the gestures collected in this study.

We used the \$N\$ multi-stroke recognizer [1] to recognize gestures. We analyzed only gestures from the *shape* and *symbol* categories, as many gesture recognizers, including \$N\$, are designed to handle glyph gestures such as these, but not taps or directional flicks. For each of the *shape* and *symbol* gestures collected in this study, we performed a recognition test using 3 sets of gestures: (1) all other *shape* and *symbol* gestures created by that participant, (2) *shape* and *symbol* gestures created by blind participants, and (3) *shape* and *symbol* gestures created by sighted participants. For (2) and (3), the creator of the gesture being tested was excluded from the set. Recognition was considered correct if the correct gesture was the top recognition result.

Table 1 summarizes the results of this analysis. A Chi-Square test revealed that gestures from sighted participants were significantly more likely to be recognized correctly ($\chi^2_{(1,N=3624)}=56.19$, $p<.001$). Recognition accuracy also differed based on the recognizer's training set. For sighted participants, recognition accuracy was higher when tested against gestures from other sighted participants than when tested against gestures from blind participants ($\chi^2_{(1,N=1250)}=35.96$, $p<.0001$). Surprisingly, recognition accuracy was also higher for blind participants when tested against sighted gestures ($\chi^2_{(1,N=1166)}=7.33$, $p<.01$). This result seems counterintuitive, as we would expect that a blind person's gestures would be more similar to gestures from other blind people, and thus that recognition accuracy would be higher when the recognizer was trained with blind gestures. A closer examination reveals that recognition accuracy for blind participants' *symbol* gestures was much higher when tested against sighted gestures than when tested against other blind gestures. This difference may be due to blind participants' unfamiliarity with some of the *symbol* gestures. Because these gestures were less familiar to the blind participants, there may have been greater variation in how they were performed. Thus, *symbol* gestures performed by blind people may have been better matched to sighted participants' gestures because the sighted gestures were internally consistent, even if they

differed from blind participants' gestures in other ways. However, this explanation is merely speculative, and further work is needed to understand the effects of blind people's gesture performance on recognition accuracy.

	vs. Self	vs. Blind	vs. Sighted
<i>Symbol</i> gestures by blind	68.2% (46.6)	52.0% (50.0)	61.3% (48.7)
<i>Shape</i> gestures by blind	59.6% (49.4)	44.9% (50.0)	44.9% (50.0)
<i>Symbol</i> gestures by sighted	72.5% (44.7)	63.2% (48.3)	78.7% (41.0)
<i>Shape</i> gestures by sighted	68.9% (46.6)	57.8% (49.7)	73.3% (44.5)

Table 1. Each gesture was tested against (1) the participant's other gestures, (2) blind gestures, and (3) sighted gestures.

Table cells report mean and SD recognition accuracy.

DISCUSSION

The overarching goal of this research is to set future directions for the design of touch screen applications, and to promote accessible touch screen interaction techniques that work equally well for both blind and sighted people. The studies described here address two primary questions: first, given the choice, would blind people prefer to perform different gestures than sighted people? Second, do blind people perform gestures differently than sighted people even when performing the same gestures?

In response to the first question, we discovered significant differences in the gestures chosen by blind and sighted participants. Blind participants in our study showed strong preferences for gestures that used screen corners, edges, and multi-touch. Furthermore, when asked to invent gestures, blind participants in our study adopted two techniques that are rare in most current touch screen user interfaces: using a second finger or hand to begin a mode, and touching areas of the screen that correspond to keys on a QWERTY keyboard. These gestures were used by a majority of blind participants in our study, including participants who had experience with other touch screen-based devices. While it remains to be seen whether these gestures would be useful or efficient in a real world application, their popularity in this study suggests that they may have potential.

In answering the second question, we uncovered a number of performance characteristics that differentiate gestures produced by blind people from those produced by sighted people. Gestures produced by blind participants were larger, slower, and featured greater variation in size than those produced by sighted participants. Although some of these differences may seem intuitive, identifying these differences in real performance data, and quantifying them, deepens our understanding of how blind people perform gestures. We have identified several important metrics related to gesture performance by blind people, including location accuracy, form closure, and line steadiness. These

factors have a significant impact on how gestures drawn by blind people may look different than those drawn by sighted people, and to our knowledge this is the first time these features have been discussed and measured in relation to how blind people perform gestures.

A relevant outcome of the second study is that some blind participants did not know how to perform some of the gestures used in our protocol, including letters, numbers, and other symbols. Even some blind participants who knew these symbols sometimes pointed out that other blind people might not know them. This is not surprising when we consider that many blind people have never learned how to print, but it is an important and non-obvious consideration when choosing gestures for an application.

Design Guidelines for Accessible Touch Screens

Based on the results of these two studies, we offer preliminary advice on how to design future touch screen-based applications for both blind and sighted users:

Avoid symbols used in print writing. Blind users may have limited knowledge of symbols used in print writing, such as letters, numbers, or punctuation. Even when these symbols are known, users may not be used to them or may not be comfortable performing them. If performing symbol-based gestures is an important part of the user experience, the user should be trained in how to perform these gestures or should be able to choose alternative gestures.

Favor edges, corners, and other landmarks. Locating precise spots on the touch screen surface can be very difficult for a user who cannot see the screen. The physical edges and corners of a touch screen are useful landmarks for a blind person. Placing critical functions in these areas will improve accessibility and reduce the likelihood that the user will trigger these functions accidentally.

Reduce demand for location accuracy. Blind users may be less precise in targeting specific areas of the screen, including edges and corners. This problem can be reduced by increasing target size or by allowing approximate targeting methods, such as allowing a user to touch near a target and then explore with their finger to locate it more precisely.

Limit time-based gesture processing. Blind people may perform gestures at a different pace than sighted people. Thus, using the gesture's speed as a recognition feature or as a parameter (as in kinetic scrolling) may result in increased errors for blind users.

Reproduce traditional spatial layouts when possible. Objects with familiar spatial and tactile layouts, such as a QWERTY keyboard or telephone keypad, are instantly familiar to many blind people. Reproducing these layouts may make it easier for a blind person to learn and use a new interface.

FUTURE WORK

The present work has uncovered several promising directions for creating new accessible interaction techniques for touch screens. In particular, the QWERTY keyboard-like interaction technique was very popular among our blind participants, and thus merits further investigation. In addition, our participants suggested a number of possible avenues for improving text entry on a flat touch screen, including handwriting Braille characters, using the Perkins Braille chording technique, or using variations of a telephone keypad. These techniques may hold promise for both blind people and sighted people in eyes-free situations.

A second area of opportunity is in improving gesture recognition accuracy for blind people. Our results show that blind participants experienced significantly reduced gesture recognition accuracy when using a traditional recognizer even when the recognizer was trained with gestures from other blind participants. We envision several possibilities for improving gesture recognition accuracy for blind people, such as by preprocessing blind users' gestures before recognition, creating a modified gesture recognizer for blind gestures, or identifying a subset of gestures that can be recognized reliably when performed by blind people.

CONCLUSION

In this paper, we explored two issues related to the design of touch screen user interfaces for blind people. First, we examined blind and sighted participants' preferences for gesture-based commands on a tablet PC by asking them to invent their own gestures. We found that blind participants did in fact suggest gestures that were different than those suggested by sighted people. Blind participants favored gestures that occurred on the edges of the screen, and suggested new gestures that utilized spatial layouts that were familiar to them. Second, we examined differences in how blind and sighted people perform the same set of gestures, and presented metrics for describing how gestures produced by blind people differ from gestures produced by sighted people.

As touch screens are now one of the most common ways of interacting with computers, it is not only important that blind people can access touch screens, but also that they can do so effectively and efficiently. The present work provides new information about how blind people currently think about interacting with a touch screen, and about how they perform gestures on a touch screen as compared to a sighted person. We believe that this work will bring us closer to the creation of robust and usable touch screen interfaces that work equally well for blind and sighted people.

ACKNOWLEDGEMENTS

We thank Nicole Torcolini for her valuable feedback. This work was supported in part by National Science Foundation grant IIS-0811063. 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. Anthony, L. and Wobbrock, J.O. A lightweight multistroke recognizer for user interface prototypes. *Proc. GI '10*, Canadian Information Processing Society (2010), 245-252.
2. Baudel, T. and Beaudouin-Lafon, M. Charade: Remote control of objects using free-hand gestures. *Communications of the ACM* 36, 7 (1993), 28-35.
3. Baudisch, P. and Chu, G. Back-of-device interaction allows creating very small touch devices. *Proc. CHI '09*, ACM (2009), 1923-1932.
4. Bonner, M., Brudvik, J., Abowd, G. and Edwards, W.K. No-Look Notes: Accessible eyes-free multi-touch text entry. *Proc. Pervasive '10*, Springer (2010), 409-427.
5. Buxton, W., Foulds, R., Rosen, M., Scadden, L. and Shein, F. Human interface design and the handicapped user. *SIGCHI Bulletin* 17, 4 (1986), 291-297.
6. Buxton, W., Hill, R. and Rowley, P. Issues and techniques in touch-sensitive tablet input. *Proc. SIGGRAPH '85*, ACM (1985), 215-224.
7. Crew, S. Touch-screen gadgets alienate blind. *Reuters*, January 8, 2009.
8. Crossan, A. and Brewster, S. Multimodal trajectory playback for teaching shape information and trajectories to visually impaired computer users. *ACM TACCESS* 1, 2 (2008), 1-34.
9. Epps, J., Lichman, S. and Wu, M. A study of hand shape use in tabletop gesture interaction. *CHI '06 Extended Abstracts*, ACM (2006), 748-753.
10. Goldreich, D. and Kanics, I.M. Tactile acuity is enhanced in blindness. *Journal of Neuroscience* 23, 8 (2003), 3439-3445.
11. Guerreiro, T., Lagoa, P., Nicolau, H., Goncalves, D. and Jorge, J. From tapping to touching: Making touch screens accessible to blind users. *IEEE Multimedia* 15, 4 (2008), 48-50.
12. Guimbretière, F., Stone, M. and Winograd, T. Fluid interaction with high-resolution wall-size displays. *Proc. UIST '01*, ACM (2001), 21-30.
13. Heller, M.A., Wilson, K., Steffen, H., Yoneyama, K. and Brackett, D.D. Superior haptic perceptual selectivity in late-blind and very-low-vision subjects. *Perception* 32, 4 (2003), 499-511.
14. Kamel, H.M. and Landay, J.A. A study of blind drawing practice: Creating graphical information without the visual channel. *Proc. ASSETS '00*, ACM (2000), 34-41.
15. Kane, S.K., Bigham, J.P. and Wobbrock, J.O. Slide Rule: Making mobile touch screens accessible to blind people using multi-touch interaction techniques. *Proc. ASSETS '08*, ACM (2008), 73-80.
16. Landau, S. and Wells, L. Merging tactile sensory input and audio data by means of the Talking Tactile Tablet. *Proc. Eurohaptics '03*, (2003), 414-418.
17. Law, C. and Vanderheiden, G. The development of a simple, low cost set of universal access features for electronic devices. *Proc. CUU '00*, ACM (2000), 118.
18. Liu, J., Pinelle, D., Sallam, S., Subramanian, S. and Gutwin, C. TNT: Improved rotation and translation on digital tables. *Proc. GI '06*, Canadian Information Processing Society (2006), 25-32.
19. McGookin, D., Brewster, S. and Jiang, W. Investigating touchscreen accessibility for people with visual impairments. *Proc. NordiCHI '08*, ACM (2008), 298-307.
20. Morris, M.R., Wobbrock, J.O. and Wilson, A.D. Understanding users' preferences for surface gestures. *Proc. GI '10*, Canadian Information Processing Society (2010), 261-268.
21. Sadato, N., Pascual-Leone, A., Grafman, J., Ibanez, V., Deiber, M.P., Dold, G. and Hallett, M. Activation of the primary visual cortex by Braille reading in blind subjects. *Nature* 380, 6574 (1996), 526-528.
22. Sánchez, J. and Aguayo, F. Mobile messenger for the blind. *Proc. ECRIM '06*, Springer (2006), 369-385.
23. Sánchez, J. and Maureira, E. Subway mobility assistance tools for blind users. *Proc. ECRIM '06*, Springer (2006), 386-404.
24. Stöbel, C., Wandke, H. and Blessing, L. Gestural interfaces for elderly users: Help or hindrance? *Gesture in Embodied Communication and Human-Computer Interaction*, Springer (2010), 269-280.
25. Tinwala, H. and MacKenzie, I.S. Eyes-free text entry with error correction on touchscreen mobile devices. *Proc. NordiCHI '10*, ACM (2010), 511-520.
26. Van Boven, R., Hamilton, R., Kauffman, T., Keenan, J. and Pascual-Leone, A. Tactile spatial resolution in blind Braille readers. *Neurology*, 54 (2000), 2230-2236.
27. Vanderheiden, G.C. Use of audio-haptic interface techniques to allow nonvisual access to touchscreen appliances. *Proc. HFES* 40, (1996), 1266.
28. Wobbrock, J.O., Morris, M.R. and Wilson, A.D. User-defined gestures for surface computing. *Proc. CHI '09*, ACM (2009), 1083-1092.
29. Yfantidis, G. and Evreinov, G. Adaptive blind interaction technique for touchscreens. *Universal Access in the Information Society*, 4 (2006), 328-337.
30. Zhai, S. and Kristensson, P. Shorthand writing on stylus keyboard. *Proc. CHI '03*, ACM (2003), 97-104.
31. Zhao, S., Dragicevic, P., Chignell, M., Balakrishnan, R. and Baudisch, P. EarPod: Eyes-free menu selection using touch input and reactive audio feedback. *Proc. CHI '07*, ACM (2007), 1395-1404.