

“With most of it being pictures now, I rarely use it”: Understanding Twitter’s Evolving Accessibility to Blind Users

Meredith Ringel Morris¹, Annuska Perkins², Catherine Yao^{1,3},
Sina Bahram⁴, Jeffrey P. Bigham⁵, Shaun K. Kane⁶

¹Microsoft Research; ²The University of Washington; ³Massachusetts Institute of Technology;
⁴North Carolina State University; ⁵Carnegie Mellon University; ⁶University of Colorado Boulder
merrie@microsoft.com, annuska@uw.edu, catyao@mit.edu,
sbahram@ncsu.edu, jbigham@cmu.edu, shaun.kane@colorado.edu

ABSTRACT

Social media is an increasingly important part of modern life. We investigate the use of and usability of Twitter by blind users, via a combination of surveys of blind Twitter users, large-scale analysis of tweets from and Twitter profiles of blind and sighted users, and analysis of tweets containing embedded imagery. While Twitter has traditionally been thought of as the most accessible social media platform for blind users, Twitter’s increasing integration of image content and users’ diverse uses for images have presented emergent accessibility challenges. Our findings illuminate the importance of the ability to use social media for people who are blind, while also highlighting the many challenges such media currently present this user base, including difficulty in creating profiles, in awareness of available features and settings, in controlling revelations of one’s disability status, and in dealing with the increasing pervasiveness of image-based content. We propose changes that Twitter and other social platforms should make to promote fuller access to users with visual impairments.

Author Keywords

Social media; Twitter; blindness; accessibility.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI).

INTRODUCTION

Worldwide, 39 million people are blind [32]. Blind computer users typically use screen reader software that converts on-screen information to aural speech or refreshable Braille; screen readers require accessible content in order to work. In particular, images require an alternative text description (called “alt text,” since web page authors can specify such text in the “alt” field of an image’s HTML tag). On web pages, alt text is often absent or insufficiently detailed [3, 15,

24]. Automated image description generation (e.g., [2, 4, 5, 14, 19, 30]) is not robust enough to be useful in practice.

Social media have become an increasingly important part of modern life. For example, more than 302 million people use Twitter every month, sending more than 500 million tweets per day [28]. Many people who are blind use social media to stay connected with friends, family, colleagues, and people with shared interests [5, 34]. However, with the ease of taking and sharing photos afforded by smartphones, social media are becoming increasingly image-based. For instance, Facebook users uploaded over 350 million photos per day in 2013 [25]. Twitter was created in 2006, and initially had a very simple text-based interface (consequently making it popular among people who are blind [5]), but began hosting embedded images in tweets on June 1, 2011 [11].

Although some blind people take and share images [18], the increasing pervasiveness of imagery lacking alt text is problematic for this constituency (to our knowledge, no major social media platforms support alt text). The rising image-heaviness of social media impedes blind users’ ability to engage fully with this increasingly important medium, as explained by one blind participant in our study:

“In an increasingly visual culture, it’s vital that people with visual impairments have equal access to information on social media... one of the most frustrating experiences I encounter is sitting down after a long day of work to scroll through my tweets and have a few laughs with my friends, only to realize that an article I want to read or a picture or video a friend has tweeted is inaccessible.”

Our research aims to pave the way toward making social media more accessible to people who are blind. In this paper, we focus on Twitter since prior studies have indicated that Twitter is disproportionately popular among blind people due to its history as a simple, text-based medium [5] and since the majority of tweets are publically accessible for analysis due to the medium’s default privacy conventions. We address the following research questions:

RQ1: What are blind users’ goals in engaging with Twitter? What barriers exist to achieving these goals, and how might technology be designed to facilitate them?

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.
CHI’16, May 07 - 12, 2016, San Jose, CA, USA
Copyright is held by the author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3362-7/16/05...\$15.00
DOI: <http://dx.doi.org/10.1145/2858036.2858116>.

RQ2: How are people who are blind currently using Twitter? How does their use differ from sighted users? Do these differences reflect different goals, or do they reflect accessibility problems? Do usage differences between blind and sighted users represent possible privacy leakages about disability status?

RQ3: Is Twitter becoming less accessible to people who are blind over time? What types of inaccessible content are prevalent, and what types of technologies might be appropriate to mitigate accessibility challenges associated with different Twitter behaviors and content types?

RELATED WORK

Social media use is on the rise; more than 70% of online American adults were using social media as of September 2014 [12]. Researchers have begun to examine the extent to which social technologies are used and usable by people impacted by disabilities, such as parents of children with special needs [1, 23] and other serious health concerns including postpartum depression [10] and eating disorders [7]. Our focus is on social media use and usability for people who are blind. Three prior studies have looked specifically at how people who are blind use social media [5, 31, 34]. Brady et al.'s research [5] provided the first insights into this area; their survey results established that people who are blind do indeed use social media, and identified that Twitter adoption is higher than typical for this user group due to Twitter's (historically) relatively simple, text-based interface. Wu and Adamic [34] conducted a log-based study of the activity patterns of Facebook users who were accessing the service via a screen reader; unlike [34], we supplement logs with survey data that provides subjective and qualitative feedback from blind social media users, and we model traits differentiating blind and sighted users. Voykinska et al. [31] also studied blind users of Facebook via interviews and surveys that offered insights into the strategies blind Facebook users adopt to work around the accessibility challenges of interpreting the image-heavy site. Buzzi et al. [6] report a case study of a blind user's difficult interactions with Twitter. Our research builds upon these findings – like [31] and [34], we are providing deep insight into how people who are blind interact with a specific social technology (in this case, Twitter, a choice of technology inspired by Brady et al.'s findings regarding its importance to this constituency; Wu and Adamic [34] noted that their findings may not extend to other media such as Twitter, so our work complements prior work to provide a richer picture of how findings generalize or differ across media).

Twitter is a microblogging social media service on which users post messages (“tweets”) limited to 140 characters. Since June 1, 2011, visual content (images and video) can be embedded directly into a tweet to supplement the 140 characters of text [11]. Twitter was identified by Pew as among the most popular social media platforms, being used by 23% of online American adults in September 2014 [12]. Twitter's dual role as a social network and a news source

makes it particularly important for many users [22]. There are several Twitter client applications designed specifically for screen reader users, such as EasyChirp [13], Chicken Nugget [8], and TheQube [26]. These apps often provide special functionality, such as compatibility with refreshable Braille displays; keeping abreast of evolving functionality and use trends is a challenge for third-party clients.

The increased pervasiveness of smartphones with high-quality cameras has made it simple for people to capture and share imagery. This trend is reflected in the quantity of images being shared in social media – several hundred million photos are uploaded to Facebook every day, for instance [25]. This paper is the first to present data demonstrating the growing proportion of and importance of images embedded in Twitter posts.

In addition to posting photos of daily life or news events, people also use imagery in tweets in informal ways that further confound accessibility for blind users by taking formerly accessible content (text) and rendering it as imagery that cannot be interpreted by a screen reader. For example, “screenshorting” [17, 33] is a trend in which users post photos or screen captures of blocks of text as images embedded in a tweet, thereby circumventing Twitter's 140 character text limit. Many users also circumvent the official “retweet” functionality by instead posting a photo or screen capture of an original tweet as an embedded image; these unofficial retweets may have a variety of motives, such as avoiding implicitly “liking” a tweet by not adding to its official retweet count, circumventing the 140 character limit by giving the retweeter more room to add commentary, and/or making it more difficult for government regimes to track who is supporting various ideologies [27].

Web protocols have long supported and encouraged site designers to add “alt” tags to their HTML code in order to provide verbal descriptions for images – this text is read aloud by a screen reader when a blind user visits a compliant web page; however, there is no affordance in Twitter (or any other major social media that we are aware of) for adding alt text to images embedded in posts. One grassroots effort to retrofit alt text into tweets is the Alt Text Bot [9]; if a Twitter user forwards a tweet with an embedded image to the Alt Text Bot's account, then the service uses the Cloudsight API (an undisclosed mix of human-powered and computer-vision-powered captioning), and tweets back a short caption. Some Twitter users may choose to use the 140 characters of their tweet to describe an embedded image; our findings reveal that this is quite rare, with only 11% of multimedia tweets having text that might serve as a useful description.

Many researchers are working to develop automated captioning technologies. Semi-automated captioning systems (i.e., human-in-the-loop labeling provided via crowd-powered systems [2, 4, 5, 30]) or fully-automated systems (i.e., approaches combining computer vision, machine learning, and natural language generation technologies [14, 19]) currently exist, but are not practical

for automatically captioning social media imagery for blind users due to issues such as cost, privacy, and response time for human-in-the-loop systems, and issues such as accuracy and generalizability beyond limited training image types for ML-based approaches. Our findings contribute information about the variety and types of imagery being shared on Twitter that can inform the design of enhanced automated approaches for captioning such imagery for blind users.

METHODS

We created an online survey to gather the perspectives of current Twitter users who are blind on their needs and concerns regarding this form of social media. The survey also collected the Twitter handles of the respondents, which we used to gather account features from the Twitter “firehose,” a data source containing all tweets from public accounts. We also gathered the same firehose data for a control set of users to compare the characteristics of blind user’s accounts with those of sighted users.

Survey Method

We used SurveyGizmo to design an online survey, and pilot-tested the survey with screen reader users to help ensure accessibility. The survey was available from March 4th to March 15th, 2015. It had 26 questions, and took a median of 12 minutes to complete.

Because the survey was about Twitter use, we recruited participants via Twitter by sending tweets advertising the survey from our personal and organizational accounts; these tweets used hashtags that people in our target demographic are likely to attend to such as *#a11y* (a general hashtag related to accessibility) and *#csun* (a hashtag related to a conference on accessible technologies that occurred during the survey period). We also purchased advertising on Twitter to promote our tweets. Twitter does not offer blindness as a demographic for ad targeting; instead, we targeted our ads to users who had ever tweeted with the hashtag *#a11y* or who followed accounts related to blindness, such as accounts associated with prominent advocacy organizations (e.g., AFB, NFB) or with services used by people who are blind (e.g., screen reader companies, guide dog organizations). To incentivize participation, we contributed \$1 to participants’ choice of the AFB, NFB, Benetech, or Lighthouse International for each completed survey.

The survey’s consent form indicated that participants must be at least 18 years old, have a Twitter account, and be blind (as defined by needing to use a screen reader to interact with computers or mobile devices). We also informed participants that their Twitter profile page and public tweets would be analyzed as part of the study. To facilitate accessibility and privacy, all questions not relating to qualifying criteria were not required, so some participants left some items blank if they did not feel comfortable disclosing them.

Survey Participants

We received 117 completed surveys and 20 that were partially complete. We manually checked the validity of all

Table 1. Motivations for using Twitter among 112 blind users.

Motivation to use Twitter	%
Become informed of news/current events	88.4%
Enjoy entertainment and/or humor	72.3%
Learn about research/issues related to blindness	70.5%
Socialize with friends or family	63.4%
Participate in advocacy related to blindness	58.0%
Facilitate career/professional networking	52.7%
Meet new people	50.9%
Meet new people who are also blind	43.8%

137 Twitter handles provided by these respondents by viewing the associated profile page in a web browser. Five handles did not resolve to valid profile pages; those respondents’ data were removed from analysis, leaving 132 users (112 with completed surveys, 20 partially complete).

All 132 respondents used screen readers (most commonly VoiceOver, JAWS, and NVDA). Respondents’ ages ranged from 18 – 67 years old, with a median age of 35 and a mean of 37.8 years. 70.5% identified as male and 19.7% as female (9.8% didn’t specify). Respondents were from a diverse set of 19 countries; the most common were the U.S. (47.7%), U.K. (13.6%), Canada (11.4%), and Australia (4.5%).

All 132 respondents’ Twitter handles were used for our analysis of account properties, but we only analyzed the answers from the 112 fully-completed surveys.

Firehose Analysis

We used the Twitter firehose to analyze public account and tweet properties of the 132 Twitter handles provided by our survey participants. We analyzed data from the six months preceding our survey (Sept. 1, 2014 through Feb. 28, 2015).

In order to compare the properties and activities of blind users of Twitter to sighted users, we assembled a control set of an additional 132 accounts. To create the control set, we used the Twitter firehose to randomly sample all distinct Twitter handles that sent at least one public, English tweet on February 28, 2015. We manually verified that each account belonged to a non-celebrity, non-spammer person (and not to an organization or business) by viewing the associated profile page. We also made sure the profile page contained no indications the user was blind, since our intent was for the control sample to represent sighted Twitter users.

Of the 132 valid Twitter handles provided by our blind survey participants, 16 had to be eliminated from firehose analysis – 13 of these were protected accounts whose tweets were not available to us, and 3 of these were accounts that sent no tweets during the six-month data collection period (presumably those account owners use their accounts primarily to read others’ tweets). This left us with 116 accounts known to belong to people who are blind. We reduced the size of the control set to 116, to match. Thus, our firehose data consisted of six months of tweets and profile properties from 232 public Twitter accounts, 116 belonging to our blind survey takers and 116 from our control sample.




Figure 1. This sample Twitter profile page uses the default profile image (an egg) and the default header image (a solid blue rectangle). The user’s Twitter handle appears below the profile image, followed by the bio.

FINDINGS

In this section, we present the findings from the survey, firehose analysis, and manual inspection of profile pages.

Motivations for Using Twitter

We asked respondents why they use Twitter, allowing them to select options from a list (Table 1), as well as a write-in option. 8.0% provided write-in responses, such as “to state opinions,” “to discover new music,” and “to document my daily life.” 84.8% of respondents identified with at least one motivation for using Twitter that was related to blindness (rows 3, 5, and 8 in Table 1).

Visual Imagery in Tweets

We asked how often respondents asked others on Twitter to describe the contents of images in tweets. Choices were *frequently (nearly every day)*; *occasionally (a couple times a week)*, *rarely (a couple times a month)*, or *never*. The majority (55.4%) reported never asking other Twitter users to describe imagery. 28.6% reported doing this rarely, 12.5% occasionally, and 3.6% frequently. We also asked whether respondents intentionally follow users who retweet images and add descriptions to them; 20.5% reported doing so. One respondent noted not only following such users, but actively engaging with them, “In the past, I’ve reached out to people who use good image descriptions and thank them...” Another noted that finding users who can give good descriptions is challenging, “I’m hesitant to ask for photo descriptions because people don’t know how to give them...”

We asked respondents what information they want about the visual content of tweets. A few people indicated this would not interest them (e.g., “None. I would not be interested in looking at visual content on Twitter at all”; “None – if people tweet inaccessibly, I stop following them.”). However, the vast majority expressed interest in having alt text for images (e.g., “tags or descriptions”; “short alternative description”; “what it is, color, and what’s in the background”). Many noted nuance that would be appreciated in captions (“if there is a visual joke, then to understand why it is humorous”; “clear description of: colours, shapes, patterns, lighting, focus, facial expression, action, body language, clothing, location, whether image contains text and what it says”; “Anything the author of said visual media wishes to highlight... I might as well know what the author intends to

display when tweeting visual media.”). Some noted that videos often posed less of a challenge than still images, due to the audio track providing some additional context (e.g., “Videos, as long as they have audio, don’t bother me”).

Profile Page Characteristics

Figure 1 shows an example Twitter profile page, illustrating the placement and defaults for visual components. A user’s *Twitter handle* (username) and *profile image* appear adjacent to every tweet they send; clicking on a particular Twitter handle or searching for a particular user takes one to that user’s profile page, where the *header image* and optional 160 character *bio* and location fields are also visible. If the account is public, all of that user’s tweets are also shown on the profile page in reverse chronological order.

We manually inspected the profile pages associated with all 132 Twitter handles of blind users collected through our survey. We performed the same coding process on the 132 control accounts (before each the sets were reduced to 116 to accommodate protected tweets – all 132 accounts in each set had public profile pages that could be examined).

Each profile page was coded on the following dimensions:

- Did the bio field explicitly state that the user was blind, allude to blindness (e.g., by mentioning interests in accessible technology, guide dogs, screen readers, Braille, etc.), or provide no clues as to the user’s vision abilities?
- Was the profile image the default (egg)?
- Was the header image the default (solid rectangle)?

We examined the blind users’ accounts’ bio fields to see whether they disclosed information about vision status. 20.5% had blank bios. 19.7% explicitly mentioned visual impairment (e.g., “Went blind in Nov. 2012”; “chocoholic who happens to be blind”; “Blind sportswriter”). 25.0% alluded to blindness by mentioning related concepts (e.g., “Seeing Eye grad not currently working a dog”; “Braille transcriber by trade”). The remaining 34.9% had bios that did not provide any clues regarding the user’s vision status.

We asked participants if their primary Twitter profile image was the default (Figure 1), or whether they had customized it. 22.3% reported that their image was the default, 50.0% reported having customized it, and 27.7% reported being unsure what their profile image was. We then compared this self-report to our manual coding of these users’ profile images. For those who reported being unsure of whether they had a non-default profile image, most (23 of 31) did indeed have the default (including in this number two people who had non-egg but completely blank profile images). Of the 55 who reported having customized their image, only 85.5% were correct (8 actually did have the default egg). Of the 25 people who reported having the default profile image, 76.0% were correct, but 6 actually did have a customized image. Some survey respondents didn’t even realize there was a default; one participant noted, “Just learned of the ‘egg’ terminology [from the description provided in the survey].”

For the respondents who said that they had customized their profile image, we asked how they selected the image. We analyzed the responses only from the 47 people who truly had a non-default image. The methods of choosing a profile image fell into three categories – help from a sighted friend, using a known (pre-existing, labeled) image, and taking or creating a new image oneself. Four responses were ambiguous and couldn't be placed with certainty into any of these three categories (e.g., “uploaded a photo”). 51.1% asked a sighted friend for assistance in creating a profile image (e.g., “My wife took the photo of me on my iPhone and said it looks good”; “I am a Star Trek fan, so I had a friend who can see choose a good [Star Trek] image for me”). 29.8% were cases where the person used a known photo (e.g., “I used a portrait taken for work purposes”). 12.7% were images created specifically for their Twitter profile by the blind user (e.g., “I took the picture with my phone”).

We asked those who indicated that they had left the image as the default to explain why. Some said that imagery is not important to them (e.g., “I don't care about pictures”). Others mentioned privacy concerns (e.g., “Because I do not want my picture displayed on social media or on the internet”). A few indicated not taking time (e.g., “I just never bothered...”). However, the most common response theme related to accessibility issues, either not being aware the profile existed or being unable to make the changes easily (e.g., “I've not been able to change it on the Twitter page with my screen reader”; “didn't know you could change it, or how to do it”).

We also asked whether respondents had customized their profile's “header image” (Figure 1). 17.0% reported being unsure, 8.0% reported that they had customized it, and 75.0% reported they had not. As with profile images, our participants were sometimes incorrect about the content of their header images – five who reported that they did not customize their header image actually had, and five who reported that they had changed it actually had the default.

Of the 132 control accounts, only 1.5% had not customized their profile image, and only 16.7% had not customized their header image. A Pearson chi-square test shows that the proportion of blind vs. control users with a custom profile image is significantly different (57.6% of blind users vs. 98.5% of control users): $\chi^2(1, N=264) = 64.4, p < .001$. A Pearson chi-square test also shows that the proportion of blind vs. control users who have customized their header image is significantly different (11.4% of blind users vs. 83.3% of control users): $\chi^2(1, N=264) = 137.1, p < .001$.

Based on the firehose analysis of the 116 blind and 116 control accounts' properties, 23.3% of control users had empty bios in their profile and 19.8% of blind users did. This is not a statistically significant difference. We also analyzed the length in characters of the bio field. Blind users had significantly longer bios (92.9 characters vs. 51.6 for control users), $t(230) = -5.685, p < .001$.

Account Characteristics

We measured users' account age in days as of March 1, 2015. For the control users, this was 1103.97 days, and for blind users it was 1749.88 days. This is a significant difference, $t(230) = -7.816, p < .001$, suggesting that blind users were more likely to be relatively early adopters of Twitter, a finding consistent with prior self-report studies on the social media habits of people who are blind [5] (an alternative possibility may be that more experienced blind Twitter users than typical self-selected to participate in our survey).

We also examined the number of followers the accounts in each set had, normalizing by the account age in days (since accounts accumulate followers over time). The normalized follower count for control users is .732, while blind users have a much lower normalized follower count of .261, which is a significant difference, $t(230) = 2.601, p = .01$.

Examining the number of friends, also normalizing by account age, control users had .896 normalized friends, whereas blind users had .763; this was not a statistically significant difference. There was also no significant difference in friend-to-follower ratio.

The final account property we examined was whether or not the account had geolocation enabled; if geolocation is enabled, a precise latitude/longitude is attached to each tweet when the tweet is made from a device (such as a smartphone) that is able to provide this data. 44.8% of control users had geo-location enabled on their account, and 56.0% of blind users had geolocation enabled. This difference is marginally significant according to a Pearson Chi-Square test, $\chi^2(1, N=232) = 2.914, p = .088$.

Tweet Characteristics

Normalizing tweet volume by account age provides a rate of tweeting (mean tweets per day). For control users this was 7.77, for blind users it was 9.12. This was not a statistically significant difference ($t(230) = -.754, p = .452$).

Our data indicated whether a tweet contained embedded multimedia, such as a photo or video (note that this refers only to multimedia embedded directly within a tweet, not to URLs within tweets that may point to media hosted elsewhere). Blind users are much less likely to share embedded visual content on Twitter than the control users – 23.43% of control users' tweets contained embedded multimedia, compared to 4.78% of blind users' tweets, which is a significant difference: $t(230) = 10.990, p < .001$.

We also examined what percent of tweets were “original” tweets as opposed to retweets. Blind users are much less likely to retweet messages than control users, with 27.33% of blind users' tweets being retweets, compared to 34.66% of control users' tweets, $t(230) = 2.257, p = .025$.

We analyzed the most frequently used hashtags among the blind and control user accounts. All words in which a “#” character preceded a sequence of letters and/or numerals were parsed out as hashtags. Any punctuation appearing at

Table 2. The most popular hashtags used by the 116 blind and 116 control group Twitter users over a six-month period.

Hashtag	# users	# tweets	category
Blind Users			
ally	73	3014	disability
blind	60	320	disability
fb	57	3808	tech
accessibility	56	1575	disability
audio	42	1307	tech
Control Users			
wcw	22	62	meme
tbt	17	31	meme
ferguson	15	98	news
neverforget	15	39	news
relationshipgoals	13	43	meme

the end of terms was removed, and case was normalized to lowercase. We then counted the number of tweets from our data set containing each hashtag, as well as how many unique users from our set ever used that hashtag.

Table 2 shows the most popular hashtags for each group. The blind users have much more convergence (higher numbers of unique users) on common hashtags, which may reflect shared interests of this group and common purpose in their use of Twitter (i.e., to find and/or share information related to blindness). The most popular hashtags for the blind users related to disability/accessibility and/or to technologies (such as “audio”), whereas the most popular hashtags for the control group appear to be memes (“wcv” = “women crush Wednesday”; “tbt” = “throwback Thursday”) and terms related to trending news events, suggesting differential uses of Twitter by these two groups (advocacy and seeking and sharing information related to accessibility and technology for the blind users, contrasted with entertainment, socializing, and news consumption for the sighted users). Note that there may be some biases toward the use of certain hashtags among our blind user set as an artifact of our recruitment method (some of our recruitment was accomplished through tweets and twitter ads targeted toward disability-related hashtags such as #ally).

Predictive Power

Our findings indicate that there are statistically reliable differences in the properties and behavior associated with the Twitter accounts of people who are blind versus people who are sighted. This suggests that there may be privacy leakage regarding peoples’ vision status – while a subset of the participants in our study chose to explicitly indicate in their Twitter profile that they were blind, most did not. However, our findings indicate that implicit cues that users may not be aware of could nonetheless reveal their disability status; this privacy leakage carries with it a risk of discrimination, such as online harassment or differential advertisement targeting.

To explore the extent to which automatic detection of visual impairment status may be possible, we performed a binomial logistic regression analysis on the 116 blind users’ accounts

Table 3. Binomial logistic regression model to differentiate blind and sighted Twitter users. $R^2 = .805$.

Variable	B	Exp(B)	S.E.	Sig.
Bio length in chars	.013	1.013	.006	.033
Bio is blank	1.138	.871	.976	.888
Account age in days	.001	1.001	.001	.072
Normalized # friends	.449	1.567	.589	.446
Normalized # followers	-.655	.519	.729	.369
Friend/follower ratio	-.026	.974	.067	.693
Mean tweets per day	.007	1.007	.019	.727
Percent photo tweets	-10.336	.000	2.892	.000
Percent RTs	.023	1.023	.013	.077
Geolocation enabled	1.168	3.216	.559	.037
Default profile photo	3.000	20.085	1.300	.021
Default bkgd image	2.596	13.405	.630	.000

and 116 control users’ accounts for which we had complete data. This model correctly classified 94.0% of control users as being sighted and 92.2% of the blind users as being blind. The Nagelkerke R Square value for this model is .805 (i.e., the model accounts for 80.5% of the variance between the blind and control accounts).

The most significant predictor variables in this model were the percent of tweets originating from the account that contained multimedia ($p < .001$), whether the account had the default background image on the profile page ($p < .001$), whether the account had the default “egg” as the profile photo ($p = .021$), the length of the account’s bio ($p = .033$), and whether the account had geolocation enabled ($p = .037$). Table 3 shows the details of the model. Note that we did not include linguistic characteristics of tweets (such as hashtags) as features in our model, but the commonalities in hashtags used by many people who are blind suggests that content-related features may also represent a similar privacy leakage.

A conservative rule of thumb to avoid overfitting is to have at least ten examples per factor [29], which we do. However, to further address possible concerns about overfitting, we re-ran the data with ten-fold cross validation using Weka; overall accuracy was 90.5%, with true positive rates of 92.2% for control and 88.8% for blind classes.

Prevalence of multimedia tweets

To analyze whether the amount of multimedia content in tweets is increasing over time, we sampled all public English language tweets sent between noon and 2 p.m. UTC on the first day of each month for the eighteen-month period from January 2014 through June 2015, a total of 159,163,658 tweets. We examined the tweets’ metadata to determine if they contained multimedia. Both embedded photos and embedded videos count as multimedia and are not differentiated in the metadata available to us since videos are embedded as a “photo thumbnail” of the first video frame. For the subset of tweets made from February 2015 onwards, we have additional metadata that allows us to distinguish whether embedded multimedia is a video thumbnail versus a static image; from this data, we find that video is still

relatively rare –93.6% of embedded multimedia in tweets consists of static imagery, while 6.4% are videos. The proportion of tweets containing embedded multimedia has been rising. In January 2014, 15.4% of tweets contained multimedia, rising to 28.4% by June 2015 (Figure 2).

Separating out tweets as “original” tweets (the first time a tweet is sent) versus retweets, we find that multimedia was embedded in 13.2% of original tweets in January 2014, rising to 23.2% by June 2015. However, retweets exhibit a far higher percentage of embedded multimedia, rising from 22.8% of retweets in January 2014 to 42.2% in June 2015. A paired samples t-test comparing the percent of multimedia original tweets versus retweets each month shows a significant difference, $t(17) = 31.81, p < .001$. This suggests that tweets containing imagery are among the most “important” according to popularity and interest levels.

To better understand the nature of imagery being shared on Twitter, we randomly sampled 900 English tweets issued on June 1, 2015 that contained embedded imagery. We coded these tweets and images with regards to three questions:

- What is the type of image embedded in the tweet?
- How important is the embedded image to interpreting the meaning of the tweet?
- Would the text of the accompanying tweet serve as a useful description for the image?

Using an open-coding methodology, two researchers iteratively developed coding schemes for these questions by

Table 4. Frequency of image types embedded in 900 English-language image-containing tweets from June 1, 2015.

Image Type	Frequency
Photographs	64.4%
Images with Embedded Text (e.g., ads)	11.5%
Pictures of Text	9.0%
Screenshots	7.0%
Drawings	5.4%
Graphs	1.1%
Inspirational Quotes	0.9%
Unofficial Retweets	0.7%
Memes	0.0%

viewing and labeling a separate random sample of embedded-imagery tweets. To illustrate the robustness of the coding scheme, these two researchers then redundantly coded 200 tweets from the set of 900 (inter-rater reliability measures for each question are reported as we discuss the findings for each, below); a single researcher then labeled the remaining 700 tweets. Because of the size of images, the firehose feed does not store the actual embedded image, but rather a URL pointing to the image; about 16% of these URLs were broken links, likely because either Twitter or the user deleted the tweet after posting. Tweets containing broken image links were discarded from analysis, leaving 756 of the original 900 tweets (and leaving 172 of the 200 in the set used for computing IRR measures).

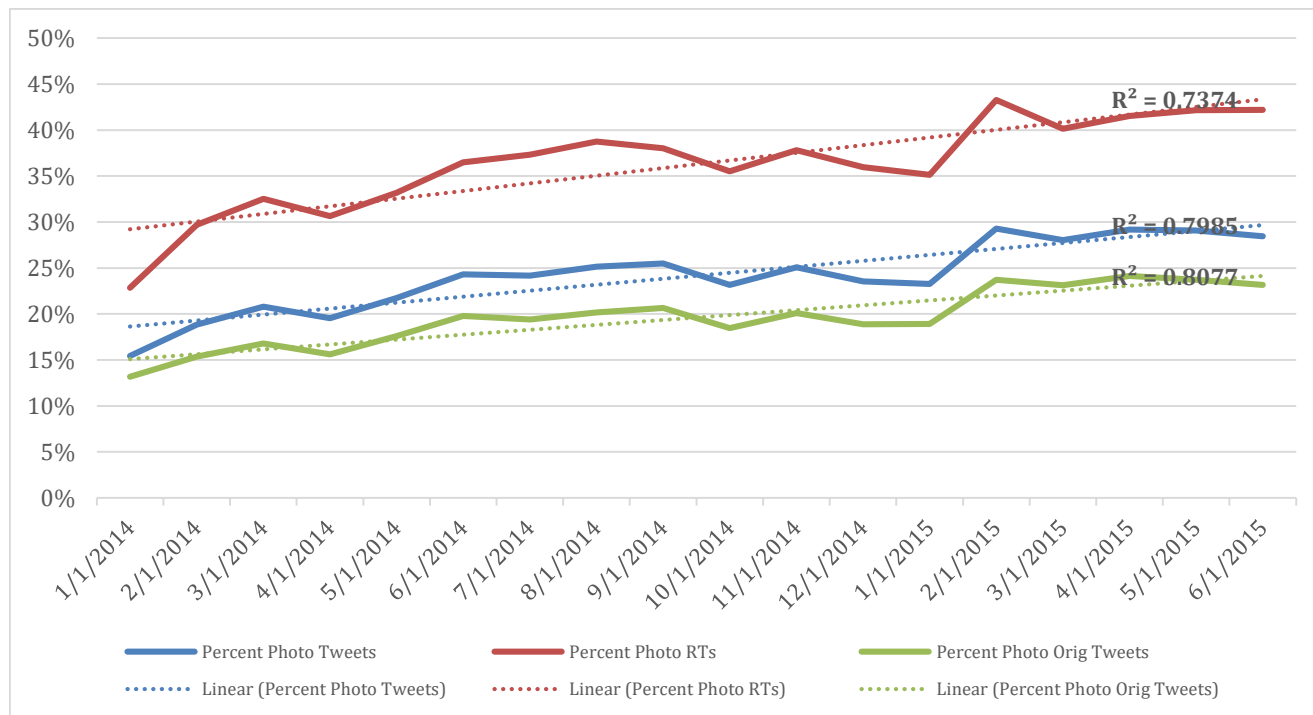


Figure 2. The percent of English-language tweets that contain embedded multimedia is increasing over time. The dashed lines show best-fit line and associated R^2 values for linear regression. Blue represents the percentage of all tweets that contain multimedia, green represents only original tweets, and red represents retweets.




Figure 3. Nine types of imagery commonly embedded in tweets. Top row: photograph; drawing; picture of text (“screenshot”); image with embedded text; screenshot. Bottom row: graph; inspirational quote; meme, unofficial retweet.

What is the type of image embedded in the tweet?
 Determining the types of images embedded in tweets facilitates identifying what captioning solutions might be most appropriate. Our coding scheme identified nine categories of imagery appearing in tweets (Figure 3). To illustrate the robustness of the scheme, we calculated the Cohen’s Kappa for the 172 images coded by two raters; the Kappa score was 0.8595, indicating substantial agreement (particularly given the large number of categories). Table 4 shows the distribution of image types in our sample of tweets. When people think of “images” in tweets, the default assumption is typically of photos; while photos comprised the largest category of imagery (64%), the diversity of image types on Twitter suggests that a variety of different captioning approaches may need to be merged to adequately cover the types of imagery shared in social media.

How important is the embedded image to interpreting the meaning of the tweet?

Answering the question of how important the embedded image is to the point of the tweet provides insight into whether the abundance of embedded imagery significantly degrades blind users’ Twitter experience; if the majority of images add no significant value to their tweets, then captioning solutions may not be necessary. Our coding scheme used an ordinal scale with the codes being *very important* (a user hearing the text of the tweet without viewing the image would not be able to understand the tweet’s meaning), *somewhat important* (a user hearing the text of the tweet without viewing the image would be able to partially interpret the tweet’s meaning), and *nonessential* (a user hearing the text of the tweet without viewing the image would be able to fully appreciate the point of the tweet). To measure inter-rater reliability, we calculated the Spearman’s correlation coefficient between the two researchers’ answers to this question, finding $r = 0.53$, an indication of moderate to strong agreement. Overall, 55.6% of the 756 embedded images were rated as very important to interpreting the

meaning of the tweet, 19.4% as somewhat important, and 25.0% as nonessential. This indicates that embedded imagery are generally integral to interpreting their respective tweets, and thus that accessible descriptions of this imagery is important for blind users’ Twitter experience.

Would the text of the accompanying tweet serve as a useful description for the image?

We wanted to know if tweets themselves provide text that could serve as a caption for blind users, so we coded how well tweet text described embedded images’ contents. Using a three-point scale, we coded the tweet’s utility as a caption as either *good* (a blind user hearing the tweet would have a detailed understanding of the accompanying image’s contents), *minimally acceptable* (a blind user hearing the tweet would have a rough idea of the accompanying image’s contents), or *poor* (a blind user hearing the tweet would have no insight into the accompanying image’s contents). We again used Spearman’s correlation coefficient as a measure of inter-rater reliability, with $r = 0.60$, reflecting moderate to strong agreement. Of the 756 tweet-image pairs labeled, 61.8% of the tweets would be poor captions. 27.0% of the tweets would have made a minimally acceptable caption, and the remaining 11.2% would have been good captions. These findings indicate that, for the great majority of tweets, tweet text would not adequately serve as a description of the embedded imagery for people who are blind.

DISCUSSION

Our survey of blind Twitter users combined with our analysis of blind and sighted users’ public Twitter profiles and tweets and our analysis of the volume and nature of embedded imagery in public tweets allow us to reflect on our research questions, and to suggest design guidelines for improving the accessibility of social media for people who are blind.

RQ1: What are blind users’ goals in engaging with Twitter? What barriers exist to achieving these goals, and how might technology be designed to facilitate them?

Our survey of 132 blind Twitter users revealed that this user group uses social media for many of the same reasons as sighted users – for keeping up with news, for entertainment, for socializing with existing ties, and for professional networking. Additionally, the majority (85%) of our survey respondents also reported using Twitter for reasons specifically related to their visual impairment, including participating in advocacy on behalf of people who are blind, meeting new contacts who are also blind, and learning about issues and news related to blindness. Our analysis of the hashtags most prevalent among blind Twitter users provides further evidence of this specialized use of social media, reflecting the popularity of terms relating to disability and accessible technologies among this user group.

Some user goals (such as meeting new contacts who are also blind) could be facilitated more automatically if Twitter were to automatically identify users who are blind, either by searching for explicit self-identification within the user’s profile bio (45% of the blind participants in our survey had bio fields revealing or strongly hinting at their vision status) or by identifying people with a high probability of being blind through more implicit means, which our regression analysis suggests is possible. Automatic or opt-in identification might also facilitate other goals of Twitter use, such as by attempting to use automated captioning technology to label tweets with embedded images in blind users’ feeds, and/or to remind the people with blind followers that they should use images judiciously or label them clearly. However, anything involving self-identification or automatic identification of blind users raise serious privacy concerns.

Other approaches to mitigating the barriers blind users currently face in achieving their goals on Twitter are to provide data about all accounts, rather than only those associated with people who are blind. For example, in addition to listing stats such as the number of friends and followers any user has, a user’s profile page might also list stats such as what percent of their tweets contain embedded images; this information could help a blind user decide whether or not to follow a particular account. Community-generated ratings (such as ratings about whether users’ tweets provide good descriptions of embedded imagery) might also be appropriate to include in profiles, particularly for certain types of user accounts (i.e., accounts that might be expected to uphold a higher standard for accessibility than an average user, such as corporate accounts, news organizations, and/or verified celebrity accounts).

Providing the ability to add alt text for screen readers to images embedded in tweets (which would not count against a post’s character limit) would be a valuable platform-level enhancement for Twitter and other popular social media platforms. Alt tag compliance rates by typical users is likely to be low, so another important platform-level enhancement might be to facilitate letting motivated users supplement the accessibility of others’ posts (as happens with other forms of user-generated content, such as Wikipedia [21]). For

instance, platforms could pay employees or crowd workers to add in any missing alt text to embedded imagery in highly retweeted tweets, or could provide facilities to allow motivated volunteers who use the medium to add alt text to a tweet originating with another user. Providing a way that blind users can request a description of a particular image, such as by forwarding the image to a special account or marking a flag indicating the need for an automated or community-generated caption would also be a valuable addition to Twitter and other social media platforms.

RQ2: *How are people who are blind currently using Twitter? How does their use differ from sighted users? Do these differences reflect different goals, or do they reflect accessibility problems? Do usage differences between blind and sighted users represent possible privacy leakages about disability status?*

Our analysis of profiles and tweets associated with 116 blind and 116 control users indicated several differences in Twitter use. Many of these differences seem to reflect usability and accessibility challenges. For example, blind users were significantly less likely to complete their profile page (profile picture and header image) than sighted users, and many reported confusion over whether or not they had updated these fields and/or surprise that these fields existed, indicating accessibility issues with completing profiles. More blind users than control users had geolocation features enabled in their account; most users prefer to obscure their precise location for privacy reasons [16], leading us to infer the blind users may have this option enabled because the status of privacy settings and/or the ability to modify them are both not easily accessible to screen readers.

Our analysis of the most common hashtags among each user group also revealed differences. The prevalence of hashtags related to accessibility issues and technology is likely due to some of the unique social media usage motivations that our blind survey respondents described. However, the absence of certain types of hashtags, such as popular memes, from the list of top tags may reflect the usability challenges of an increasingly image-centric medium. For example, among the most common hashtags used by the control group were tags such as *#tbt* and *#wcv* that are typically associated with posting embedded imagery (for “tbt” aka “throwback Thursday” people typically post a nostalgic image, and for “wcv” aka “women crush Wednesday” people typically post images of women they admire or find attractive). This differential hashtag use, along with our finding that tweets with images are more likely to be retweeted than those not containing imagery, suggests that many popular aspects of Twitter are not currently accessible to blind users.

Our regression analysis indicates that many blind users may be able to be automatically identified as blind, even if they are like the 55% of our sample that didn’t explicitly mention topics related to visual impairment in their user bio; implicit indicators such as profile page characteristics, retweeting rates, and tweet content provide a reliable basis for inferring

vision status. This may represent a privacy leakage [20] that could have negative implications for these users – for example, the platform could identify their vision status to advertisers, resulting in discrimination.

RQ3: *Is Twitter becoming less accessible to people who are blind over time? What types of inaccessible content are prevalent, and what types of technologies might be appropriate to mitigate accessibility challenges associated with different Twitter behaviors and content types?*

Our analysis of a large sample of tweets over an eighteen month period indicates that Twitter is quickly morphing from a primarily text-based medium to a primarily multimedia one. As of June 2015, 28.4% of tweets contained embedded multimedia (not including tweets containing URLs that link to external sources of imagery or video). We found that image-containing tweets are among the most popular, being retweeted at a far higher rate than text-only tweets. The prevalence of images, and the lack of a mechanism for providing alt text for embedded images, means that Twitter, once a preferred medium for people who are blind [5], is becoming increasingly inaccessible.

Our analysis of the types of embedded imagery found on Twitter reveals that images do not simply refer to photographs, but to a wide variety of visual content types. This finding suggests that automated solutions for captioning these images may not be straightforward, and may require systems that combine different technologies (computer vision for recognizing items in photos, OCR for recognizing text in “screenshots,” etc.). Some types of embedded media might be best dealt with by platform changes or social engineering approaches rather than attempts at automated or manual captioning. For example, the rise of the “screenshorting” trend [17, 33] suggests that increasing the character limit may be desirable to end users and would reduce the inaccessible images used to circumvent current limits. The use of embedded screenshots of tweets as an unofficial retweet mechanism is another example of users taking something that could be screen reader accessible text and instead sharing it as inaccessible images – Twitter may wish to consider addressing the situations that cause people to circumvent the official retweet mechanisms (which may include things like a desire to circumvent the 140 character limit in order to add long commentary to a retweet, the desire to avoid “liking” or boosting the stats of a retweet, and/or the desire to avoid being officially linked to the sender of a tweet in regimes that monitor political tweets [27]).

Limitations

Readers should bear in mind that the blind people who chose to participate in our study may have been somehow non-representative of other blind Twitter users who chose not to participate. Although we used an accessible survey template and pilot-tested with screen reader users, feedback from some participants indicated that some screen reader versions may have had compatibility issues, which may have biased our survey sample in ways that are difficult to quantify.

While rich, our Twitter firehose data does not contain information about who a given user follows; this prevented us from doing some analyses that may have provided interesting perspectives, such as identifying what percent of each blind user’s feed actually contained imagery (i.e., providing insight into whether blind Twitter users tend to follow people who tweet less imagery than is typical).

Also, note that this research implicitly assumes that our sample of control users is sighted (we manually screened their bios for any implications that they were blind and did not find any); however, it is possible (though statistically unlikely) that some members of the control set were blind.

Our study does not investigate the accessibility of Twitter to people with other disabilities; however, improved text equivalents for imagery may benefit other populations, such as other users of screen readers (people with low vision) and users with cognitive challenges (who may gain deeper understanding of content with a clear text description).

While our research focused on Twitter, we believe many of our findings have implications for other social media. Our work joins prior work [5, 31, 34] in providing evidence that people who are blind are interested in using social media; the lack of alt text capabilities on major social platforms should be a top design concern for such services. Our findings on the variety of imagery types shared socially provide new challenges that social platforms and computer vision researchers must consider when thinking of the breadth of imagery to service with emerging captioning tools. This work also illustrates how a platform that was previously quite accessible to the blind can evolve over time to become increasingly inaccessible due to a combination of new features introduced and changing end-user behaviors; this finding should be a reminder to all platforms that accessibility may not be a one-time solution, but must evolve along with a platform and its users.

CONCLUSION

In this paper, we used a multi-method approach to explore how and why people who are blind use Twitter, and the challenges they face in doing so. Our findings revealed that blind users have many goals in common with sighted social media users, but also use social media for reasons related specifically to advocacy for, education about, and socializing amongst others who are blind. We also demonstrated that profile page characteristics and tweet patterns provide reliable, albeit unintentional, means of distinguishing blind from sighted users. Our findings also showed that Twitter is becoming more image-heavy over time, and that image-based tweets are diverse, largely inaccessible, and not clearly amenable to any existing automated captioning methods. We provided several suggestions for solutions that might allow Twitter to continue to be accessible to people with visual impairments; ensuring the accessibility of social media is increasingly important given the prominent role of such new media in social and political life.

REFERENCES

1. Ammari, T., Morris, M.R., and Schoenebeck, S.Y. Accessing Social Support and Overcoming Judgment on Social Media among Parents of Children with Special Needs. *Proceedings of ICWSM 2014*.
2. Be My Eyes. <http://bemyeyes.org>.
3. Bigham, J.P., Kaminsky, R.S., Ladner, R.E., Danielsson, O.M., and Hempton, G.L. WebInSight: Making Web Images Accessible. *Proceedings of ASSETS 2006*.
4. Bigham, J.P., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R.C., Miller, R., Tatarowicz, A., White, B., White, S., and Yeh, T. VizWiz: Nearly Real-Time Answers to Visual Questions. *Proceedings of UIST 2010*.
5. Brady, E., Zhong, Y., Morris, M.R., and Bigham, J. Investigating the Appropriateness of Social Network Question Asking as a Resource for Blind Users. *Proceedings of CSCW 2013*.
6. Buzzi, M.C., Buzzi, M., and Loporini, B. Web 2.0: Twitter and the blind. *Proceedings of CHIItaly 2011*.
7. Chancellor, S., Pater, J., Clear, T., Gilbert, E., and De Choudhury, M. #thyghgapp: Instagram Content Moderation and Lexical Variation in Pro-Eating Disorder Communities. *Proceedings of CSCW 2016 (to appear)*.
8. Chicken Nugget. https://q-continuum.net/chicken_nugget/.
9. Cundiff, C. "Alt Text Bot." *AT&T NYU Connect Ability Challenge*, June 23, 2015. <http://connectability.challengepost.com/submissions/37785-alt-text-bot>
10. De Choudhury, M., Counts, S., and Horvitz, E. Predicting Postpartum Changes in Emotion and Behavior via Social Media. *Proceedings of CHI 2013*.
11. Dorsey, J. "search + photos." *The Official Twitter Blog*, June 1, 2011.
12. Duggan, M., Ellison, N., Lampe, C., Lenhart, A., and Madden, M. Social Media Update 2014. *Pew Internet and American Life Project*, January 9, 2015.
13. EasyChirp. <http://easychirp.com>.
14. Fang, H., Gupta, S., Iandola, F., Srivastava, R.K., Deng, L., Dollár, P., Gao, J., He, X., Mitchell, M., Platt, J.C., Zitnick, C.L., and Zweig, G. From Captions to Visual Concepts and Back. *Proceedings of CVPR 2015*.
15. Goodwin, M., Susar, D., Nietzio, A., Snaprud, M., and Jensen, C.S. 2011. Global web accessibility analysis of national government portals and ministry web sites. *Journal of Information Technology & Politics*, 8(1), 41-67.
16. Hecht, B., Hong, L., Suh, B., and Chi, E.H. Tweets from Justin Bieber's Heart: The Dynamics of the "Location" Field in User Profiles. *Proceedings of CHI 2011*.
17. Honan, M. The Rise of the Screenshort. *BuzzFeed News*. December 19, 2014.
18. Jayant, C., Ji, H., White, S., and Bigham, J.P. Supporting Blind Photography. *Proceedings of ASSETS 2011*.
19. Karpathy, A. and Fei-Fei, L. Deep Visual-Semantic Alignments for Generating Image Descriptions. *Proceedings of CVPR 2015*.
20. Kosinski, M., Stillwell, D., and Graepel, T. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, Vol. 110 No. 15, Feb. 12, 2013, pp. 5802-5805.
21. Kuksenok, K., Brooks, M., and Mankoff, J. Accessible Online Content Creation by End Users. *Proceedings of CHI 2013*.
22. Kwak, H., Lee, C., Park, H., and Moon, S. What is Twitter, a social network or a news media? *Proceedings of WWW 2010*.
23. Morris, M.R. Social Networking Site User by Mothers of Young Children. *Proceedings of CSCW 2014*.
24. Shi, Y. E-Government Web Site Accessibility in Australia and China: A Longitudinal Study. *Social Science Computer Review*, Fall 2006. 24: 378-385.
25. Smith, C. "Facebook Users are Uploading 350 Million New Photos Each Day." *Business Insider*, Sept. 18, 2013.
26. TheQube. <http://theqube.oire.org>
27. Tufekci, Z. Big Questions for Social Media Big Data: Representativeness, Validity and Other Methodological Pitfalls. *Proceedings of ICWSM 2014*.
28. Twitter. "Twitter Usage/Company Facts." <https://about.twitter.com/company>, retrieved July 27, 2015.
29. Vittinghoff, E. and McCulloch, C.E. Relaxing the Rule of Ten Events per Variable in Logistic and Cox Regression. *American Journal of Epidemiology*, March 2007, 165(6): 710-8.
30. Von Ahn, L. and Dabbish, L. Labeling Images with a Computer Game. *Proceedings of CHI 2004*.
31. Voykinska, V., Azenkot, S., Wu, S., and Leshed, G. How Blind People Interact with Visual Content on Social Networking Services. *Proceedings of CSCW 2016 (to appear)*.
32. World Health Organization. "Visual Impairment and Blindness" factsheet, August 2014. <http://www.who.int/mediacentre/factsheets/fs282/en/>

33. Williams, O. R.I.P. Blogging, Killed by Screenshots.
The Next Web News. April 14, 2015.

34. Wu, S. and Adamic, L. Visually Impaired Users on an
Online Social Network. *Proceedings of CHI 2014*.