Visually Impaired Users on an Online Social Network

Shaomei Wu
Facebook Inc.
Menlo Park, CA
shaomei@fb.com

Lada Adamic
Facebook Inc.
Menlo Park, CA
ladamic@fb.com

ABSTRACT
In this paper we present the first large-scale empirical study of how visually impaired people use online social networks, specifically Facebook. We identify a sample of 50K visually impaired users, and study the activities they perform, the content they produce, and the friendship networks they build on Facebook. We find that visually impaired users participate on Facebook (e.g. status updates, comments, likes) as much as the general population, and receive more feedback (i.e., comments and likes) on average on their content. By analyzing the content produced by visually impaired users, we find that they share their experience and issues related to vision impairment. We also identify distinctive patterns in their language and technology use. We also show that, compared to other users, visually impaired users have smaller social networks, but such differences have decreased over time. Our findings have implications for improving the utility and usability of online social networks for visually impaired users.

Author Keywords
visually impaired users; vision disability; social media; social networking sites; Facebook;

ACM Classification Keywords
H.5.m. Human Factors: Measurement

INTRODUCTION
Vision-impairment is a prevalent health problem worldwide: recent statistics[1] [23] show that there are 285 million visually impaired people globally and 6.6 million in the US. However, due to the lack of systematic methods to identify this particular population, there is little research of visually impaired users on the Internet. Among the few existing studies, most have focused on usability testing for improving the accessibility of specific online services or applications, with results collected through survey and/or in-person interviews [22, 12, 5, 4]. As a result, we have only a rudimentary understanding of how visually impaired people use the Internet today.

This paper is the first attempt to answer this question quantitatively with big data, with a special interest placed on how visually impaired users engage with online social networks, in this case Facebook. Along with other online social networks, Facebook has grown to become a major part of many people's online experience. 69% of American Internet users are also Facebook users, and the average time they spend on social networking sites has almost tripled since 2006[2], accounting for 18% of the total time spent online [9].

Visually impaired people engage with online social networks just as everyone else does. With technologies such as screen reader software, OCR readers, and the WAI-ARIA standard [19], visually impaired users can navigate social networking sites through desktop computers or mobile devices [22]. In a recent study of 191 blind people recruited online, 92% of the respondents reported using at least one social networking site, and 80% of them use Facebook [5]. Despite the high penetration rate of Facebook among visually impaired people, our knowledge about how they engage with Facebook is limited. There are many basic questions to be answered. What do they do on Facebook? What content do they share? Whom do they interact with? What do their social networks look like?

In this paper, we present insights about the use of Facebook by 50K visually impaired users in several perspectives, including their Facebook activities, the content they produce, and the structural characteristics of their friendship networks. Our study is motivated by the following research questions.

RQ1: Does the behavior of visually impaired users differ significantly from the behavioral patterns of other users?

There are many barriers that might prevent visually impaired users from fully engaging with Facebook. Technologically, the use of JavaScript to create highly dynamic web pages can cause problems for screen readers, and bugs related to accessibility can be harder to capture and reproduce. Practically, certain Facebook features, such as photo sharing, seem tailored to sighted users. And these make up some of the most popular features on Facebook.

On the other hand, the advance of assistive technologies such as screen readers and voice input might already have removed barriers, enabling visually impaired users the same access as anyone. For example, user studies have found that visually impaired users can complete most tasks with increasing ease through Facebook's mobile interface [22]. There are also several new photography applications that help visually impaired users identify objects by taking photos [12, 21]. One of the

[1] `https://nfb.org/factsaboutblindnessintheus`

[2] `http://en.wikipedia.org/wiki/Facebook_statistics`

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 ACM 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'14, April 26–May 1, 2014, Toronto, Canada.
Copyright © 2014 ACM ISBN/14/04...$15.00.
DOI string from ACM form confirmation

most famous users on Instagram, Tommy Edison[3], is blind since birth and has uploaded over 300 photos to tens of thousands of followers.

If visually impaired users' ability to engage with Facebook is not constrained by physical difficulties, do they interact with friends on Facebook the same way others do? In some sense, they are a minority group, and they may have special concerns about exposing themselves in a virtual environment. For example, a recent qualitative study showed that blind users feel less comfortable asking vision-related questions on Facebook than offline, due to a combination of low response rates and the concern about appearing to be overly dependent to their Facebook friends [5].

Indeed, it is not clear how other users perceive and interact with visually impaired users on Facebook. Are they aware of, or sensitive to the existence of visually impaired users around them? For example, can they tell the difference between photos uploaded by visually impaired users and photos by sighted users? If they do, will they respond to them differently? In the Facebook activity section, we not only show the difference between activities performed by visually impaired users and average users, but also the difference between the feedback received by these two groups. Then, we will explore the motif of such difference by further investigating the content produced by these two groups in the content analysis section.

RQ2: What kind of content do they produce and share on social networking sites?

Previous research [5] has shown that the visually impaired users are able to leverage social networks to ask visual questions of their friends, but hesitate to do so because of the high perceived social costs. Other research [7] found that people with cognitive disabilities, such as autism, have serious privacy concerns when talking about their conditions online.

However, engaging with online social networks has also been found to be generally positive and supportive for people with disabilities or medical conditions [7, 17]. Taken together, minority groups on online social networks have to decide how much to share about their conditions to gain social support without paying a high social cost. Are visually impaired users on Facebook facing the same dilemma? Do they share their problems and issues, or do they prefer more general topics, keeping their visual challenges behind their computer screens? When engaging with large, general-purpose social networks like Facebook, both behaviors seem reasonable. By analyzing the textual content of visually impaired users' status updates and photo captions, we will offer insights on the overall trend at a high level (in the content analysis section).

RQ3: How are visually impaired users' social networks structured?

Previous research suggested that blind users have smaller and denser social networks [5], but this result was from survey responses of a small group of blind people. To what extent is this true for visually impaired users at a large scale on Facebook? Also, homophily theory would predict that people with vision impairment are more likely to befriend each other; can we quantitatively verify this effect here? If the homophily effect is significant and visually impaired users are well-connected to each other, would that expand their social networks beyond their offline social circles, thus situate them in larger social networks?

Answers to these questions will not only offer us a broader view of how visually impaired people engage in online social networks, but could also lead to the development of more general and accurate methods for identifying visually impaired users online. In fact, some websites or products have been working on an enhanced, or even completely separate experience for visually impaired users. For example, Facebook and Twitter have developed keyboard shortcuts for common actions and site navigation, and Amazon has an accessible version with far fewer images, simpler DOM structure, and no flash or javascript[4]. Most existing work on web accessibility is still in the preliminary stage where websites strive to function correctly with accessibility technologies. Through this paper, we aim to offer deeper understanding on the needs of visually impaired users and call for fundamental design changes that suit their needs. We will discuss the design implications of our findings in the discussion section.

RELATED WORK

The value of online social networks has been widely recognized. By connecting people and propagating information among them, online social networks foster communities [1], spread new opinions and behaviors [16, 24], deliver critical information rapidly in reaction to crises [11], and even lead to societal changes [20]. For individuals, the value of online social networks varies across social groups. Many previous studies looked at how different social groups engage with online social networks. boyd et al. [3] studied the use of online social networks by teenagers and young adults, claiming that social networks are crucial for young adults to "work out identity and status, make sense of cultural cues, and negotiate public life". As shown in a study by Pew Research Center, 43% of American Internet users older than 65 are using online social networks today (mostly, Facebook), and the main function of social media for seniors is to maintain ties with family, particularly those who live far way [10]. Burke et al. studied how children and parents communicate on Facebook [6], while Morris et al. found significant changes in new mothers' site choice, post content, and post frequency pre and after child-birth [14]. The general theme found in these studies is that people of different social groups are adopting and embracing online social networks for distinct reasons, which affect the ways in which these different social groups interact with social media.

Among this large body of research on social media usage by specific social groups, there have been relatively few studies of people with disabilities. Burke et al. [7] looked at how people with autism use social media, and found that although online communication can be less stressful and more supportive, autism patients find it difficult to maintain online

[3] `http://instagram.com/blindfilmcritic`

[4] www.amazon.com/access

connections and trust online "friends". Tsaousides et al. [17] surveyed individuals with traumatic brain injury, finding that most of them either already use Facebook regularly or are interested in learning to use it. However, their further engagement with Facebook is hindered by security concerns and cognitive deficits. Indeed, compared to other social groups, people with disabilities need to overcome substantial technological and cognitive barriers to fully engage with online social networks.

Most existing work studying visually impaired people on online social networks has emphasized assessing and removing these barriers. Wentz et al. [22] evaluated the accessibility of Facebook with different interfaces by running usability studies of 15 blind people. They found that Facebook's mobile interface is more accessible than the desktop version. In addition, their study demonstrated the popularity of mobile technology (especially smart phones) among blind people. Bigham et al. [2] designed *VizWiz*, an iPhone Q&A application that allows blind people to take a photo, record a question related to the photos in audio form, and submit the question (both photo and audio) to Amazon Mechanical Turk for answers. Some recent work [4, 5] evaluated the use of VizWiz by a few thousand blind users over a year, categorized the questions they asked, and explored the appropriateness of leveraging online social networks as friend-sourced Q&A platforms for blind people. Their results are consistent with previous findings [7, 17]: while online social networks can potentially provide tremendous support to people with disabilities, the social costs of exposing one's problems and vulnerability is a serious concern for these users.

These results highlight the complexity of accessibility issues, and challenge the basic assumptions many current services and products have made when serving visual impaired users. Although developed along this line of research, our work differs from previous studies in two ways: first, we study the general patterns of how visually impaired users engage with Facebook at a unprecedented scale; second, we investigate several novel perspectives of visually impaired users' activities, for example, the textual content they generate, the interaction they perform with friends, and the structural properties of their social networks.

DATA

The analysis was done in aggregate on anonymized data on Facebook's servers. To preserve privacy, all text analysis was performed automatically. Any content cited in the paper had been posted publicly by the users.

Our analysis focuses on users who access Facebook on their iPhones through Apple's default screen reader service (VoiceOver). This selection was made because the Apple iPhone has become one of the most popular devices among blind users [4] and is equipped with free screen reader software and a variety of assistive applications (for example, TapTapSee (`http://www.taptapseeapp.com/`), VizWiz [12]). These users are able to take advantage of Facebook's mobile interface, which is more accessible and usable for blind users than the desktop version [22]. We detect a visually impaired user if he or she accesses Facebook in VoiceOver mode for at least 3 days in a month. We filter out people who turned on VoiceOver once or twice, in case they enter VoiceOver mode accidentally. Also, a sensitivity analysis with different cutoffs for VoiceOver usage (5, 7, 10) yield qualitatively similar results.

We first find all individuals who accessed Facebook from their iPhone in a one-month period between June 15, 2013 and July 15, 2013. Among them, we draw a random sample of 50K visually impaired users, and a random sample of 160K users who were active for at least 3 days in this period as the control group[5]In the rest of this paper, we refer to the first sample as the *VoiceOver sample*, and the second sample as *iOS sample*. To study the Facebook activities and content produced by the sampled users, we collected their status updates, photo uploads, comments and likes from August 4, 2013 to August 25, 2013 (in total 3 weeks), as well as all the feedback (comment and likes) on this content received within a week of posting. We picked the samples in June and analyzed their activities in August to ensure that participants in our samples are not brand new users. To compare the network structure of these two samples, we took a snapshot of all sampled users' friendship networks on July 15, 2013.

DEMOGRAPHICS

Although our study is limited to the population of iPhone users, given the large number of users we sampled, our sample consists of a demographically diverse group of visually impaired users. In this section, we will present the country, age, and gender distribution of visually impaired users in our sample, and compare them with the control group of randomly-sampled iPhone users. All the demographic dimensions we study are based on the self-reported data in user profile pages.

Country We first look at how our sampled users are geographically distributed. Table 1 shows the top 5 countries for each group, as reported in user's profile. The first thing we notice is that both populations are highly concentrated in developed countries, especially, the United States (with around one third of the users in both samples). Secondly, it is interesting to notice that the iPhone and VoiceOver seem to be more prevalent among vision-impaired users in Europe and America than Asia – Japan has the 3rd most iPhone users in the general iOS sample, but only ranks at the 14th in the VoiceOver sample. Overall, although our sample of vision-impaired users are highly skewed towards people from developed western countries, it still achieves greater geographical diversity than previous studies, containing users from 183 different countries.

Age Not surprisingly, both samples are skewed towards young adults, with over 50% of the people in their 20s and 30s, slightly less than one quarter of the sample in their teens, and around one quarter above 40 years old. There are only a few people in our sample who do not report their ages, and we also filter out people who self-report as being over 90 years old. The average age for the VoiceOver sample is 30.14 years

[5]Note that the sample sizes are not proportional to the total number of Facebook users in each group.

Rank	VoiceOver sample	iOS sample
1	US (32.5%)	US (35.5%)
2	UK (7.2%)	UK (7.3%)
3	France (4.9%)	Japan (4.9%)
4	Germany (4.3%)	Canada (3.8%)
5	Italy (4.1%)	Germany (3.6%)

Table 1: Most self-reported countries for users in VoiceOver sample and iOS sample, with percentages in each sample.

old, and for iOS sample is 30.43 years old. We do not see a significant difference in terms of age distribution between the visually impaired user group and the control group.

Gender Both genders are well represented in two samples. However, there are slightly more males (51.8%) than females (47.6%) in the VoiceOver sample, but more females (51.9%) than males in the iOS sample (47.1%).

Note that although the visually impaired users in our sample are much more diverse than previous studies in this field [22, 4, 5], it is not a fair representation of the global visually impaired population. By constraining our study to iPhone users on Facebook, we over-sample users from well-developed western countries, who are also younger and with relatively high income.

FACEBOOK ACTIVITY

As an online social network, Facebook is most commonly used to share content (e.g., status updates, photos) and interact with content shared by friends (e.g., comments, likes). We thus focus on the four most representative activities - status updates, photo sharing, comments, and likes - and study how visually impaired users engage with these four types of activities on Facebook. We try to understand:

- What do visually impaired users do on Facebook? Do they generate a different amount of content relative to other users?
- How do other users interact with the visually impaired? Are they aware of the presence of visually impaired users online? Do they engage with this population actively?

To answer these two questions, we collect all the status updates, photo uploads, comments and likes by all the users in the VoiceOver sample and the iOS sample for three weeks in August 2013, together with all the feedback (comments and likes) on this content within a week of posting. Then we compare the volume of content created and feedback received across the two groups.

In Figure 1, we break down user activity into three categories: content produced, feedback sent, and feedback received, and we show the bootstrapped and Winsorized mean for each metric over each user group along with error bars indicating the 95% bootstrapped confidence interval. For example, in Figure 1a, we count the total number of status updates and photo uploads in three weeks for each user, and plot the bootstrap mean and confidence intervals for users in iOS sample and VoiceOver sample. We can make two interesting observations from this plot: first, visually impaired users post many more

Metric	VoiceOver - iOS	p value
total photos	-9.74×10^{-5}	1.2×10^{-14}
total status updates	4.92×10^{-5}	$< 2.2 \times 10^{-16}$
total comments	2.09×10^{-5}	4.40×10^{-11}
total likes	1.83×10^{-5}	0.7572
photo comments	-8.27×10^{-6}	$< 2.2 \times 10^{-16}$
photo likes	-1.00	$< 2.2 \times 10^{-16}$
status comments received	6.49×10^{-5}	$< 2.2 \times 10^{-16}$
status likes received	2.86×10^{-5}	$< 2.2 \times 10^{-16}$
photo comments received	-7.83×10^{-6}	$< 2.2 \times 10^{-16}$
photo likes received	-4.44×10^{-5}	$< 2.2 \times 10^{-16}$

Table 2: Wilcoxon rank sum test on the difference in medians

status updates than the control group; second, although visually impaired users do upload fewer photos than users from the control group, the gap is surprisingly small - visually impaired users are producing and sharing a significant amount of photos on Facebook.

Inspired by this finding, we also separate the number of comments and likes by the type of the content being responded to: photos or status updates. As shown in Figure 1b and Figure 1c, compared to other users, visually impaired users: (1) produce more comments, but fewer comments and likes on photos posted by others; (2) receive more likes and comments on their status updates, but not on their photos.

To confirm these observations, we also compare the medians of each metric with a Wilcoxon test (since the values are not normally distributed) in Table 2. The results are qualitatively the same: VoiceOver users produce more status updates and more comments, perform fewer photo-related activities, and receive more comments and likes on their status updates but fewer photo comments and likes (all $p < 0.001$). We also notice that the differences between the two samples are quite small (except the number of likes on photos), suggesting that VoiceOver users' ability to engage with the key functions of Facebook might be confined, but is not eliminated.

Overall, VoiceOver users are highly active at generating content and giving feedback to others' content on Facebook. Moreover, the visually impaired users in our sample on average receive more feedback on their status updates (and presumably more attention) from other users on Facebook.

However, we do see evidence that VoiceOver users engage less with photo-related activities: they upload slightly fewer photos, and comment on or like fewer photos than the iOS sample users. In addition, different from the status updates, the photos uploaded by VoiceOver users do receive fewer comments or likes from others compared to the photos uploaded by users in the iOS sample.

Since some previous studies suggest that crowding-sourcing, or friend-sourcing vision-related questions on social networks can be a great resource for blind users [5, 4], we also examine the question-asking activity of visually impaired users on Facebook by looking for question marks in status updates and photo captions, a heuristic method commonly

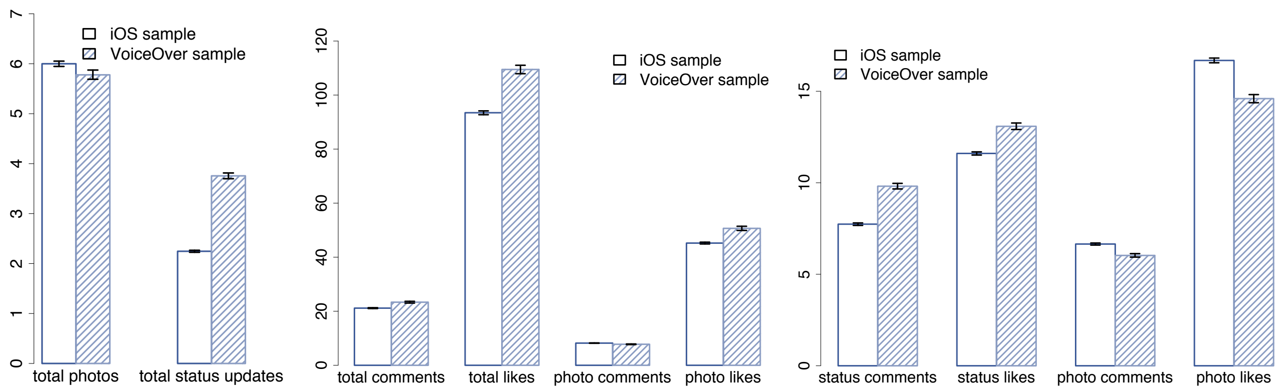

(a) Content produced (b) Feedback sent (c) Feedback received

Figure 1: Per user activity over three weeks, bootstrap averaged across user samples

used in prior studies [14, 15]. Figure 2 shows the amount of social Q&A over three weeks, including: total number of questions asked in status updates and photo captions, bootstrap averaged per capita for each sample group[6]; total number people who commented on questions posted by sampled users, again, bootstrap averaged per capita for each sample. Additionally, we calculate the ratio of questions in all status updates and in all photos (aggregated over each sample), and find: over all status updates posted by VoiceOver users, 17% of them contain a question mark, compared to 18% in the iOS sample; and for photos, only 0.4% of the photos uploaded by VoiceOver users contain question marks, exactly the same fraction as in the iOS sample. Overall, we find that the question-asking behavior is rare in both populations, and asking questions about photos is particularly uncommon (bootstrap mean is 0 for both populations). This result can partly be explained by the smaller number of photos uploaded by VoiceOver users. Also, it is consistent with previous findings that blind users are reluctant to ask vision questions of their social networks due to the high social cost perceived [5]. However, visually impaired users do ask slightly more questions than the iOS sample users, and when they ask questions through social media, they receive significantly more response than the general population does (Wilcoxon test $p < 0.0001$ for both metrics).

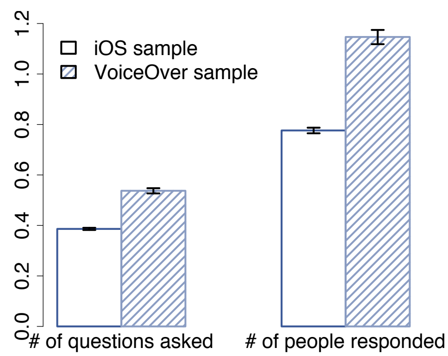

Figure 2: Social Q&A in status updates and photos

CONTENT ANALYSIS

Knowing that visually impaired users actively produce and share content on Facebook that generates feedback at higher rates than average, we would like to take a closer look at the content itself, looking for the key differences between the content shared by the visually impaired and the general population. What do visually impaired users talk about in their status updates? What kind of photos do they upload? Do they talk about their disabilities? Why do we see disproportionally more feedback on the content produced by visually impaired users?

To answer these questions, we take all the textual content in status updates and photo captions by sampled users with locale ‘en-US’ (US English), and apply the trend detection algorithm as described in [13]. Using the collection of text produced by users in the iOS sample as a baseline, we find the most representative words used by VoiceOver users with both the absolute change metric and the probability change metric. Although the probability change metric has been most recommended and widely applied in the industry[7], we include the results based on the absolute change metric because it favors words with higher frequencies [13]. Thus, it can offer a better sense about the prevalence of those selected words and prevent our results from being dominated by a small set of highly distinctive (but relatively infrequent) words.

As we have seen in the previous section, people respond more to status updates than to the photos by VoiceOver users (see Figure 1c). To better understand the difference in the nature of these two types of content, we separate the text in status

[6]Here we show the raw number of questions instead of the proportion of them in status updates, because many users did not post any status updates in the entire period thus it is not clear what the proportion means in those cases. For those who do post status updates, the median ratio of questions asked is 0 for both groups.

[7]E.g. 2010 memology: **https://blog.facebook.com/blog.php?post=466369142130**

updates from that in photo captions and run the trending term detection algorithm independently on each corpus.

At a high level, we find that visually impaired users do openly talk about issues related to vision disability and accessibility. Also, we are able to identify several technologies and applications besides VoiceOver (e.g., TapTapSee, TuneIn, Peachtree Radio) that are especially popular among visually impaired users. This finding can help us not only improve the integration of social media and these apps for a more accessible and smoother experience, but also better identify and recognize visually impaired users on the site.

In the rest of this section, we will present a more detailed analysis on status updates and photo captions.

	VoiceOver sample		iOS sample	
Rank	**Abs. change**	**Prob. change**	**Abs. change**	**Prob. change**
1	blind	blind	with	with
2	braille	braille	have	have
3	guide	sighted	this	this
4	accessible	blindness	that	that
5	sighted	goalball	just	just
6	cane	voiceover	love	love
7	audio	paratransit	time	time
8	blindness	inaccessible	like	your
9	impaired	accessible	your	it's
10	visually	impairment	what	what

Table 3: Most representative words in status updates

Status Updates

Our method for identifying the most representative words is a direct application of the two-point trends detection algorithm as described in [13]: using the status updates of iOS sample users as a training set to train a language model that predicts the frequencies of terms used in the status updates of VoiceOver sample users (testing set), we want to find the words in testing set that appear significantly more than predicted. To measure the significance of the change from "predicted" to "truth", we use normalized absolute change and the probability change metrics defined as below:

Let n_0 and n_1 denote the total number of tokens in the text from the training and testing sets respectively, and $f_0(w)$ and $f_1(w)$ denote the frequency of the word w in the training and testing sets respectively, we define the **normalized absolute change** for word w as

$$f_1(w)(n_0/n_1) - f_0(w), \tag{1}$$

and the **probability change** as

$$\binom{n_1}{f_1(w)} p_0(w)^{f_1(w)}(1 - p_0(w))^{(n_1 - f_1(w))}, \tag{2}$$

here $p_0(w) = f_0(w)/n_0$, $p_1(w) = f_1(w)/n_1$.

When adopting the algorithm to our dataset, we first filter out all terms with length less than 4 (mostly numbers and stop words). To further reduce the noise in our data, we also filter out the terms that appear fewer than 10 times in the VoiceOver sample, and fewer than 30 times in the iOS sample (since the iOS sample is 3 times bigger than the VoiceOver sample). We then normalize the language use across users by only counting each word at most once per user, which effectively reduces the algorithm's bias towards words used by users who post long, repetitive status updates. The top 10 words selected by these two metrics from all status updates in the VoiceOver user sample are shown in Table 3 (all text is converted to lowercase). To better show the difference between these two samples, we use the same algorithm but with the VoiceOver users' status updates as the baseline (training set), compute the most representative words from the status updates by iOS sample users, and show them in Table 3 as well.

As shown in Table 3, the top 10 words from the VoiceOver sample by both metrics are all related to vision disability[8], while the top words from the iOS sample are very general. Among the 12K VoiceOver users who posted any status updates, 7% of them used at least one of the 10 keywords by the abs. change metric and 5.5% used the top keywords from the prob. change metric. Meanwhile, only 0.9% and 0.4% of the iOS sample users did. Compared to other iPhone users, VoiceOver users on Facebook are much more likely to discuss vision impairment and accessibility issues found in both the physical world and on the Internet. The highly characteristic content generated by visually impaired users distinguishes them from other social media users, potentially allowing for the automatic detection of users with vision disability by their language use. Meanwhile, the uniqueness in the status updates of visually impaired users could also contribute to the higher volume of feedback from other users, as being perceived as more interesting/meaningful/important.

Photo Sharing

Knowing that visually impaired users also post photos like other iOS users, we want to assess the content of these photos, and understand whether visually impaired users upload photos that are as characteristic as their status updates. And if so, why do visually impaired users' photos not receive as much feedback from others as their status updates do?

We apply the same method as presented above, this time with the collection of text from photo captions. In Table 4, we show the top 10 most representative words that describe VoiceOver users' photo uploads (normalized by iOS sample), as well as the top 10 words used by average iOS users (normalized by VoiceOver sample).

At first glance, the keywords picked from photo captions do not appear to be as relevant to vision disability as the words from status updates are. We see some words related to listening to radio (e.g. "listening", "radio"), which makes sense since radio programs are more accessible to visually impaired users than other media such as TV and newspapers. Following this clue, we realize that most of the words shown in Table 4 are related to *specific activities or applications visually impaired users engage with* online or with their mobile

[8]Goalball is a sport developed specifically for blind athletes

	VoiceOver sample		iOS sample	
Rank	**Abs. change**	**Prob. change**	**Abs. change**	**Prob. change**
1	tunein	peachtree	this	this
2	radio	tun.in	with	with
3	listening	hatchi	love	love
4	peachtree	hatchiapp	from	gwen
5	tun.in	dailyquote	happy	happy
6	hatchi	taptapsee	have	from
7	hatchiapp	solara	just	birthday
8	dailyquote	itunes	time	have
9	taptapsee	navy/gold	birthday	little
10	bit.ly	tunein	good	just

Table 4: Most representative words in photo captions

phones. For example, *tunein* is probably from the product named "TuneIn Radio", one of the largest mobile applications for online radio (including radio stations and podcasts). And *peachtree* [9] is a popular radio station on TuneIn Radio. Photo captions containing "tunein" are mostly formulaic, such as: "I am listening to Livin' On A Prayer by Bon Jovi on WICS Radio America with TuneIn Radio **http://tun.in/seXtz**" (quote is taken from publicly shared status updates). We also find a very popular photography application for iPhone users with vision impairment - *TapTapSee*[10]. As described in the top customer review in the iTunes AppStore: "It is a camera [app] that when a picture is taken will give back a verbal description of what is seen. I use it to detect colors in order to cord Nate [sic] my wardrobe. It is one of the most helpful apps that I have on my iPhone." Hundreds of VoiceOver users in our sample have taken photos with this application and uploaded these photos to Facebook with captions like "I discovered this was a 'Nature Valley Oats N Honey Bar And Ceramic Mug On Table' with TapTapSee" and "I discovered this was a 'Coca Cola Can' with TapTapSee" (both quotes are taken from publicly shared photos on Facebook).

The top keywords from photo captions suggest that many of the photos uploaded by visually impaired users are automatically created by other apps instead of the users themselves. As a result, these photos may be viewed by others as more spammy, and thus attract less feedback than the status updates do. Meanwhile, with the popularity of photo Q&A systems such as TapTapSee and VizWiz [12], more and more blind users can get satisfactory answers to vision questions without paying the high social cost of directly polling their friends in social networks[5]. In respect of user privacy, we did not look at images directly, and thus found it difficult to separate app-generated photos from personal photos.

To summarize, we find that visually impaired users openly talk about their experiences and issues with vision disability and web accessibility. Their stories and concerns are well received and elicit active response from other users of the social network. Our trend detection algorithm is able to identify the most characteristic words and applications used by the visually impaired users, showing great potential toward a better profiling scheme for this specific population.

NETWORK STRUCTURE

In previous sections, we found that visually impaired users are actively engaging with their social networks, talking about their conditions and concerns openly, and receiving more feedback from other users. But how much of these observations can be explained by the structural properties of visually impaired users' social networks? For example, a previous study showed that users with more diverse and sparser friendship networks perform more self-censoring on Facebook [8]. Can the openness we observe in vision-impaired users be a result of their social networks being denser than average? On the other hand, the reason that visually impaired users receive more comments and likes may simply be that they have more friends (and thus a bigger audience) than an average user.

To answer these questions, we study the structural properties of the social networks around visually impaired users, focusing on the network size, density, and the interconnectivity among visually impaired users.

Network Size

Previous work has suggested that blind users have smaller social networks than average [5]. In our data, the VoiceOver median friend count of 208 is lower than the iOS median of 242 (Wilcoxon rank sum test $p < 2.2 \times 10^{-16}$). Similarly the mean and standard deviation for the VoiceOver sample is 339.9/449.8, relative to 367.5/441.0 for iOS[11].

However, the difference in network size may not have to do with users' visual impairment, but rather with the length of time they have been on Facebook (their Facebook "age"). With recent advances in web accessibility, visually impaired users are becoming able to use social networking services more easily [22]. VoiceOver users have on average been on Facebook for 38 months compared to 46 months for the iOS sample (t-test $p < 2.2 \times 10^{-16}$). To illustrate how friend count changes with Facebook age, we plot the median friend count for users who joined Facebook in the same month (see Figure 3). Our result shows that the gap in the network size of these two populations has been decreasing in recent years. Especially, when we control for Facebook age, new Facebook users (i.e., people who joined in the past 2 to 3 years) have similar network size whether they are part of the VoiceOver or iOS group.

Network Density

Another hypothesis we want to test is whether visually impaired users have more homogenous, and thus denser social networks. Previous research showed that over 95% of their blind users' Facebook social networks were comprised by friends, family and colleagues [5]. Also, as mentioned, visually impaired users engage with Facebook actively, and receive more feedback on their content than other users do. As a result, one might expect that the friends of a visually impaired user are more likely to also be friends with each other (maybe

[9] **http://tunein.com/radio/Peachtree-Radio-FM-s198932/**
[10] **http://www.taptapseeapp.com/**

[11] These numbers are much higher than previously reported [5]

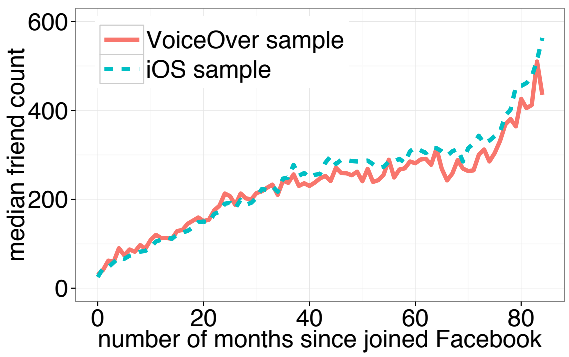

Figure 3: Network size vs. Facebook age

through both engaging with the content generated by their visually impaired friend), forming a more intimate, tightly connected social network. To measure the connectedness of a user's social network, we define the ego graph clustering coefficient of a user u_i as:

$$C_i = \frac{\text{number of edges between } u_i\text{'s friends}}{n_i \times (n_i - 1)/2} \quad (3)$$

Here n_i is the number of friends u_i has.

The greater C_i is, the denser user u_i's social network is. C_i is 0 when none of u_i's friends connect to each other, and is 1 when u_i's friends form a fully connected clique. As the clustering coefficient is in general very sensitive to the size of the ego graph (n_i) - it is much easier for small graphs to have higher clustering coefficient than big ones - we control for the size of an individuals' ego graph and plot the value of the clustering coefficient as a function of ego graph size in Figure 4. At a high level, the curves for the VoiceOver sample and the iOS sample are almost identical, showing no evidence that visually impaired users have denser social networks than the general population.

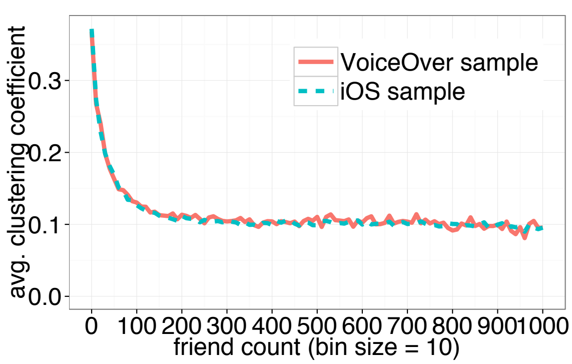

Figure 4: Ego graph clustering coefficient vs. ego graph size

We can also quantify the homogeneity of a user's social network by the number of distinct social communities among his/her friends. We use the algorithm as presented in [18] to detect and identify communities in each user's ego network, and show the distribution of community counts across users in three samples in Figure 5. As Figure 5 shows, the level of diversity in personal social networks is almost identical for users from the VoiceOver sample and users from the iOS sample, with about half of the sample having only one community, and almost 90% of the sample having no more than 3 communities. Overall, our result confirms that most visually impaired users have closely connected social networks with a few communities (presumably formed by friends, family and colleagues), but this is also true for other users on Facebook! We do not find visually impaired users to have denser than average networks over all.

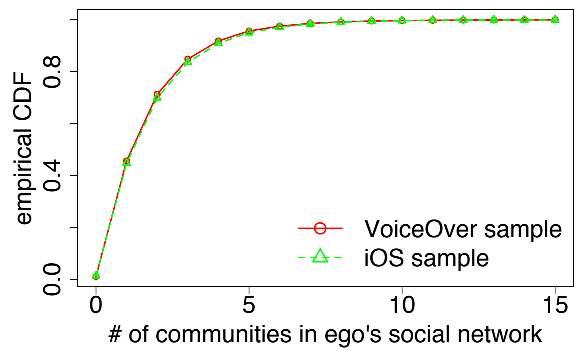

Figure 5: Distribution of number of communities

Interconnectivity among Visually Impaired Users

The last question we want to ask about the network structure of the visually impaired is whether they are more likely to be friends with other visually impaired users. Classic homophily theory would say "yes", however, given the fact that vision impaired users have relatively fewer friends, they are statistically less likely to have friends with any specific trait than people from the general iOS sample.

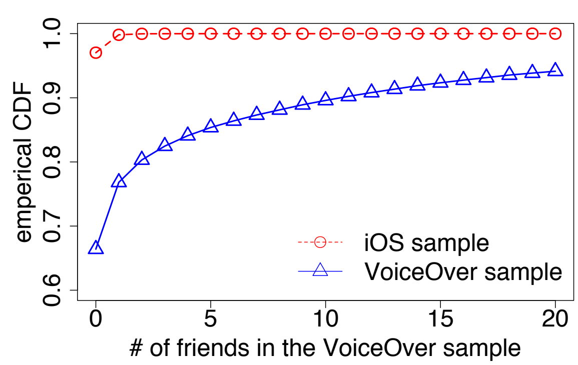

Figure 6: Distribution of friend count in VoiceOver sample

Our result supports the homophily hypothesis that visually impaired users are more likely to friend other visually impaired users. Figure 6 shows the distribution of the count of friends who are themselves VoiceOver users. Here, we can see a clear distinction between the CDF curves for these two groups: while less than 2% of users in the iOS sample who have at least one friend in the VoiceOver sample, over 20%

of the VoiceOver sample users have at least one friend who is also in the sample, and around 10% of the VoiceOver sample users have more than 10 friends using VoiceOver as well.

Such significant interconnectivity among visually impaired users would potentially introduce structural clustering of them on the Facebook network, which may eventually lead to self-organized communities of visually impaired users. This might be a factor that contributes to the characteristics of content produced and shared by visually impaired users. The interconnectivity can also be very helpful when trying to auto-detect the presence of visually impaired users.

DISCUSSION

Limitations

By analyzing the activities of visually impaired users on Facebook at an unprecedented scale, we are able to uncover high-level patterns in the behavior and language usage of visually impaired users, which might not be present or observed at a small-scale. However, there are also some limitations.

We only looked at visually impaired users who access Facebook with VoiceOver and iPhones, which is a very small subset of visually impaired people online and in the world. Although a popular device among blind people [4], the iPhone still occupies a niche market due to its price and marketing strategy as a high-end luxury phone. As a result, most of the users we studied are from developed countries with relatively high socio-economic status, even though vision impairment is more prevalent in developing countries [23]. Even within the more developed regions like North America and Europe, we cannot claim to study all visually impaired users, as a significant fraction of visually impaired people access the Internet uses non-Apple phones, desktop computers, or other screen reading software such as Microsoft Jaw[12]. Reaching out to the visually impaired population who are under-privileged and un-equipped with technologies and understanding their needs is an important area for future work.

Our data does not include important background information about the participants, such as the history and severity of their vision impairment. In many current studies, vision impairment is categorized into several levels from “low vision” to “completely blind”. People with different types of vision impairment have different technological needs and usage patterns. For example, people with low vision can use screen magnifiers if they encounter problems while using screen reader software. With a screen magnifier, they can also engage with photos the same way as others do.

We only studied Facebook users, who may behave differently on other online social media such as Twitter and LinkedIn. Previous research has shown that communication patterns on Twitter differ from those on friendship-based social networks like Facebook [24]. In the same way, we may find that our results are specific to the context of Facebook and do not generalize to other online social networks.

Privacy considerations limited the depth of our analysis. Although a content-based analysis of photos would be interesting, we did not view any images. Instead, we analyzed captions as a proxy for content.

To address these biases, and to better explain the patterns we discovered and understand what they mean to the experience of visually impaired users, we would need in-depth knowledge from qualitative studies.

Design implications

As discussed in the Content Analysis section, the characteristics of content produced by visually impaired users show potential for developing algorithms that auto-detect visually impaired users of social media. Auto-detection of users’ vision impairment could help online services to recognize this underserved market. There has been an increasing focus on visual and dynamic elements on website designs, especially, for online social networking sites. We find that, although technology has enabled visually impaired people to take photos and identify certain objects in digital images, the usage of image-related features are much lower than sighted users. Also, it is harder to navigate a webpage with dynamic elements with screen reader softwares and keyboard. In fact, a site with simple, linear structure, and minimal visual elements and dynamic content is much more user-friendly to the visually impaired community. Today, only a few websites have recognized this market and designed a separate version for visually impaired users.

Auto-detection of visually impaired users could also enable online services to better adapt to their needs. Currently, there is no a standard and effective method to identify visually impaired users. Even for websites that do have a more accessible version (e.g., `www.amazon.com/access`, `msite.facebook.com`), users need to know the address of these alternative versions and manually input these longer urls to visit them. If websites could detect the visual impairment of a visitor, they could redirect him/her to the accessible version, just like many websites offer to redirect visitors with low bandwidth or no JavaScript to the basic HTML version.

CONCLUSION

This paper studied visually impaired users’ use of Facebook. We find that they participate in all the main social activities (posting status updates, comments, and likes) just like the general Facebook population, including some photo-related features (e.g., photo comments and likes). On the other hand, when visually impaired users produce and share personal content such as status updates, they receive on average more feedback (i.e., comments and likes) from others. These findings suggest the utility of Facebook as a platform for visually impaired users to openly share their experience, voice their concerns, and receive attention and support from others.

We find a difference in size but not density of visually impaired users’ social networks relative to the rest of the Facebook users. We note that the difference in network size between these two groups has diminished over time, demonstrating progress towards an increasingly equal and accessible

[12] `http://webaim.org/projects/screenreadersurvey4/`

online environment. Also, visually impaired users are much more likely to have friends who are also visually impaired.

Finally, our analysis of the content generated by visually impaired users revealed highly characteristic keywords related to vision disability. Similarly, many photos are associated with (and perhaps auto-uploaded by) applications popular with visually impaired users (e.g. online radio, photography apps). These distinctive features of visually impaired users' online activity pave the way for developing machine-learning models to recognize visually impaired users beyond the population of iPhone/VoiceOver users.

In the future, we would like to expand our analysis to a more general population of visually impaired users online. With this initial big-data study of how visually impaired users engage with Facebook, we hope to bring more attention to the visually impaired population online and invite more research efforts toward understanding and addressing their needs.

ACKNOWLEDGMENTS

We thank Jeffrey Bigham and Cameron Marlow for insightful discussions. Also, many thanks to the accessibility team at Facebook, especially Jeff Wieland and Ramya Sethuraman. We also thank Moira Burke for feedback and editing. Lastly, we thank the reviewers for valuable feedback.

REFERENCES

1. Backstrom, L., Huttenlocher, D., Kleinberg, J., and Lan, X. Group formation in large social networks: membership, growth, and evolution. In *Proc KDD'06* (2006), 44–54.

2. Bigham, J. P., Jayant, R., Ji, H., Little, G., Miller, A., Miller, R. C., Miller, R., Tatarowicz, A., White, Y., White, S., and Yeh, T. Vizwiz: Nearly real-time answers to visual questions. In *Proceedings of UIST'10* (2010).

3. boyd, d. Why youth (heart) social network sites: The role of networked publics in teenage social life. *MacArthur Foundation Series on Digital Learning Youth, Identity, and Digital Media Volume* (2007).

4. Brady, E., Morris, M. R., Zhong, Y., White, S., and Bigham, J. P. Visual challenges in the everyday lives of blind people. In *Proc CHI'13* (2013), 2117–2126.

5. Brady, E. L., Zhong, Y., Morris, M. R., and Bigham, J. P. Investigating the appropriateness of social network question asking as a resource for blind users. In *Proc CSCW'13* (2013), 1225–1236.

6. Burke, M., Adamic, L. A., and Marciniak, K. Families on facebook. In *Proc ICWSM'13* (2013).

7. Burke, M., Kraut, R., and Williams, D. Social use of computer-mediated communication by adults on the autism spectrum. In *Proc CSCW'10* (2010), 425–434.

8. Das, S., and Kramer, A. Self-censorship on facebook. In *Proc ICWSM'13* (2013).

9. Delaney, J., Salminen, N., and Lee, E. Infographic: The Growing Impact of Social Media. `http://www.sociallyawareblog.com/2012/11/21/time-americans-spend-per-month-on-social-media-sites/`.

10. FITZGERALD, B. R. Report: 43% of seniors on web use social media. *The Wall Street Journal* (8 2013). `http://blogs.wsj.com/digits/2013/08/05/report-43-of-seniors-on-web-use-social-media`.

11. Goolsby, R. Social media as crisis platform: The future of community maps/crisis maps. *ACM Trans. Intell. Syst. Technol. 1*, 1 (Oct. 2010), 7:1–7:11.

12. Jayant, C., Ji, H., White, S., and Bigham, J. P. Supporting blind photography. In *Proc ASSETS'11* (2011), 203–210.

13. Kleinberg, J. Temporal dynamics of on-line information streams. In *Data Stream Management: Processing High-speed Data*, M. Garofalakis, J. Gehrke, and R. Rastogi, Eds., Springer (2004).

14. Morris, M. R. Social networking site use by mothers of young children. In *Proc CSCW'14* (2014).

15. Paul, S. A., Hong, L., and Chi, E. H. Is twitter a good place for asking questions? a characterization study. In *Proc ICWSM'11* (2011).

16. Romero, D. M., Meeder, B., and Kleinberg, J. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proc WWW'11* (2011), 695–704.

17. Tsaousides, T., Matsuzawa, Y., and Lebowitz, M. Familiarity and prevalence of facebook use for social networking among individuals with traumatic brain injury. *Brain Inj. 25*, 12 (2011), 1155–62.

18. Ugander, J., Karrer, B., Backstrom, L., and Marlow, C. The anatomy of the facebook social graph. *ArXiv e-prints* (Nov. 2011).

19. Web Accessibility Initiative - Accessible Rich Internet Applications. `http://en.wikipedia.org/wiki/Web_Accessibility_Initiative#Accessible_Rich_Internet_Applications_.28WAI-ARIA.29`.

20. Weber, I., Garimella, V. R. K., and Batayneh, A. Secular vs. islamist polarization in egypt on twitter. *unpulished material* (2013). `http://www.qcri.com/app/media/1900`.

21. Weiss, T. C. Taptapsee camera app for visually impaired. *Disabled World* (5 2013). `http://www.disabled-world.com/assistivedevices/apps/taptapsee.php`.

22. Wentz, B., and Lazar, J. Are separate interfaces inherently unequal?: an evaluation with blind users of the usability of two interfaces for a social networking platform. In *Proc iConference'11* (2011), 91–97.

23. World health organization: Global data on visual impairments 2010. `http://www.who.int/blindness/GLOBALDATAFINALforweb.pdf`.

24. Wu, S., Hofman, J. M., Mason, W. A., and Watts, D. J. Who says what to whom on twitter. In *Proc WWW'11* (2011), 705–714.