

“It’s almost like they’re trying to hide it”: How User-Provided Image Descriptions Have Failed to Make Twitter Accessible

Cole Gleason
Carnegie Mellon University
Pittsburgh, Pennsylvania
cgleason@cs.cmu.edu

Patrick Carrington
Carnegie Mellon University
Pittsburgh, Pennsylvania
pcarrington@cmu.edu

Cameron Cassidy
Carnegie Mellon University
Pittsburgh, Pennsylvania
ccassidy@tamu.edu

Meredith Ringel Morris
Microsoft Research
Redmond, Washington
merrie@microsoft.com

Kris M. Kitani
Carnegie Mellon University
Pittsburgh, Pennsylvania
kkitani@cs.cmu.edu

Jeffrey P. Bigham
Carnegie Mellon University
Pittsburgh, Pennsylvania
jbigham@cs.cmu.edu

ABSTRACT

To make images on Twitter and other social media platforms accessible to screen reader users, image descriptions (alternative text) need to be added that describe the information contained within the image. The lack of alternative text has been an enduring accessibility problem since the “alt” attribute was added in HTML 2.0 over 20 years ago, and the rise of user-generated content has only increased the number of images shared. As of 2016, Twitter provides users the ability to turn on a feature that allows descriptions to be added to images in their tweets, presumably in an effort to combat this accessibility problem. What has remained unknown is whether simply enabling users to provide alternative text has an impact on experienced accessibility. In this paper, we present a study of 1.09 million tweets with images, finding that only 0.1% of those tweets included descriptions. In a separate analysis of the timelines of 94 blind Twitter users, we found that these image tweets included descriptions more often. Even users with the feature turned on only write descriptions for about half of the images they tweet. To better understand why users provide alternative text descriptions (or not), we interviewed 20 Twitter users who have written image descriptions. Users did not remember to add alternative text, did not have time to add it, or did not know what to include when writing the descriptions. Our findings indicate that simply making it possible to provide image descriptions is not enough, and reveal future directions for automated tools that may support users in writing high-quality descriptions.

CCS CONCEPTS

• **Human-centered computing** → **Accessibility**;

KEYWORDS

Accessibility, Alternative Text, Image Description, Blind, Twitter

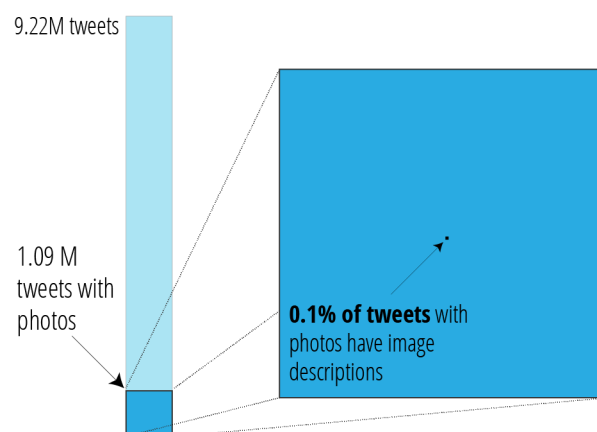


Figure 1: We sampled 9.22 million tweets, collecting 1.09 million with images. Only 0.1% of these tweets contained alternative text for people with vision impairments created using the opt-in image description feature on Twitter.

ACM Reference Format:

Cole Gleason, Patrick Carrington, Cameron Cassidy, Meredith Ringel Morris, Kris M. Kitani, and Jeffrey P. Bigham. 2019. “It’s almost like they’re trying to hide it”: How User-Provided Image Descriptions Have Failed to Make Twitter Accessible. In *Proceedings of the 2019 World Wide Web Conference (WWW ’19)*, May 13–17, 2019, San Francisco, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3308558.3313605>

1 INTRODUCTION

Twitter and other social media platforms have become increasingly visual [24] over the past decade as media such as photos, videos, and GIFs have become more prevalent as content. As visual media makes up a larger portion of total content on a social media platform, such as Twitter, such platforms risk becoming less accessible to people with vision impairments who use screen reading software to access the site [24].

An estimated 39 million people around the world are blind, and many access online sites through screen reader software. They are typical users of social networks [6, 34], but they are not always afforded access to content on social networks like Twitter [24], which are increasingly part of public discourse. Twitter is a platform for

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW ’19, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6674-8/19/05.

<https://doi.org/10.1145/3308558.3313605>

members of the media to disseminate and discuss information, as well as users to interact with celebrities in a different format than more traditional media [35]. Many research efforts have examined participation on Twitter in the context of politics around the globe [3], especially during elections [28]. It is important that all visual content on Twitter, including those relating to these topics, be accessible to people with vision impairments so they may have equal participation in public life on the Internet.

Social media platforms have taken different approaches to make visual content more accessible to blind users, although all approaches provide textual descriptions (alternative text) to user-posted images. For instance, Facebook and Instagram both automatically tag each image uploaded to the site using image detection and recognition algorithms [18, 36]. Users can edit and override this text after the image has been posted. Twitter, on the other hand, allows users to add their own descriptions when the image is posted, provided the user has previously enabled that feature.

The automatic alternative text provided by Facebook is always available, but is not yet trustworthy to blind users compared to high-quality alternative text written by humans [36]. However, few users have enabled Twitter’s image description feature for their account, and those that have do not always remember to write alternative text. We were interested in further understanding why Twitter users chose to enable and use this feature to better understand what motivates them to provide image descriptions. This understanding will help us improve similar features on social networking sites and increase the number of users providing high-quality alternative text for their social media content.

To understand the current state of alternative text provided on Twitter, we collected a sample of 1.09 million photo tweets and found that only 0.1% contained alternative text. By looking at a sample of posts with alternative text written in English, we found that 83.4% of human-written descriptions were of high quality. We then interviewed 20 Twitter users who had written image descriptions to understand their motivations for writing them.

Our findings suggest that very few users enabled the ability to post alternative text, indicating that Twitter could increase accessibility by turning the feature on by default. Those who do use the feature often do so infrequently, but generally provide alternative text of high-quality (excluding automatic posts from bots). Users who do use the feature could still benefit from training or tools that would help them write better image descriptions. We suggest that researchers or Twitter community members who wish to improve accessibility for the site develop these tools and measure their impact on the accessibility of social media content.

2 RELATED WORK

Our investigation of the accessibility of Twitter is related to prior research on both manual and automatic methods to generate image descriptions, and the usage of alternative text on the web.

2.1 What are Image Descriptions?

Image descriptions, often referred to as alternative text or “alt text”, are captions for images online or in other software. Alternative text is most commonly encountered on webpages, as it was added to the HTML 2.0 specification in 1995 [4]. However, at the time it was

intended for users who used non-graphical browsers or preferred not to render images when accessing the web. Visually impaired users are not mentioned as an intended audience until the HTML 4.0 specification in late 1997 [9]. Accessibility of images for screen-reader users is one of the most commonly cited reasons to add alternative text to images today, but it is also recommended in case the image does not load for sighted users. Image descriptions are also expected to be added to software in other domains, including mobile phone applications on the iOS [2] and Android [14] platforms. Commercial software can also add alternative text to images in documents, such as Microsoft Word and Adobe Acrobat [1, 22].

Alternative text is not solely meant for human consumption, as various search engines consume it and use it to rank pages [19]. In fact, image descriptions have been used for a number of different applications including “semantic visual search, visual intelligence in chatting robots, photo and video sharing in social media, and aid for visually impaired people to perceive surrounding visual content” [17]. Image labels, captions, and descriptions provide a solid foundation for many of these kinds of applications.

In this paper we focus on alternative text image descriptions for people with vision impairments. Thus, our evaluation revolves around the descriptive quality of captions for a human, as in [25].

2.2 Methods to Generate Alternative Text

Currently, alternative text is primarily created by the developers of the website or authors of the website content. This text is manually written, and authors are recommended to follow the Web Content Accessibility Guidelines [10]. However, as many images on the web are not labelled correctly or at all [5, 16], researchers have sought to automatically generate image descriptions. Both for online accessibility and in situ photos, subsets of individual objects in photos can be labelled using object recognition algorithms [36]. More complex images can be described by scene description models, although this is usually less robust than human labelling. A commercial example of this is Microsoft Seeing AI [21]. However, researchers should remain skeptical of automatically-generated descriptions until they are of consistently high-quality, as prior work has also illuminated that there is a threshold for when image descriptions and captions can do more harm than good [25].

Because many images are complex and difficult to describe, especially for photos with a specific intent, the best alternative text is typically provided by human labelling. As relying on the original content authors to provide accessible images has not proven fruitful, researchers have proposed multiple ways to share alternative text for images between users [12] and collaboratively make websites more accessible without permission from the owner [26]. Guinness *et al.* proposed the Caption Crawler to automatically retrieve alternative text attached to the same image elsewhere on the web through reverse image searches, which achieves a similar goal without active crowdsourcing [15].

2.3 Alternative Text on the Web

The majority of image content on the web does not contain image descriptions. Bigham *et al.* found that less than 40% of significant images on the top 500 high-traffic websites contained alternative text [5]. This motivated the authors to create a tool to generate

alternative text automatically from surrounding web context and optical character recognition. They found that on top-ranked sites by traffic, they automatically generated captions for around 50% of not-described images. In a historical analysis of websites from 1997-2002, Hackett *et al.* found that websites were getting increasing complex and less accessible [16]. Notably, in both this study and Bigham *et al.*, government websites tended to be more accessible than other groups. A more recent (2017) survey of top websites by Guinness *et al.* found 20-35% of images lacking alt text in various categories [15]. Over the last ten years, the number of accessible images on the web has not noticeably changed.

With the rise of social media platforms, a significant amount of image content on the web are now generated by end-users, not website authors. This has led to a large amount of content being inaccessible, as users did not have the option to add descriptions to their posts. Morris *et al.* found that over 25% of English tweets in June 2015 contained an image, and Twitter did not allow alternative text to be added at the time [24]. The post text itself was not a substitute for alt text, as the authors found that only 11.2% of tweets would serve as good descriptions for their accompanying images. In 2016, Twitter added an opt-in feature for users to write image descriptions for their images, which we examine in detail.

3 THE STATE OF TWITTER ACCESSIBILITY

We sought to quantify the usage of the image description feature across Twitter to understand how the introduction of image descriptions have made Twitter more accessible. We describe the image description feature, in general, followed by overall creation of alternative text by individual, popular, and public accounts.

3.1 Adding Image Descriptions

In 2016, Twitter added a feature to allow users to add image descriptions in their tweets that contain images [27]. By default this feature is disabled. The feature must be enabled by going to *Settings and privacy* -> *Accessibility* and clicking *Compose image descriptions* [30]. A note below the checkbox says “Adds the ability to describe images for the visually impaired.” and a link provides more information in Twitter’s documentation. Once the feature is enabled, a visual cue appears when a user uploads an image on the website or mobile application (Figure 2). The tweet author may add a description of up to 420 characters to each image (a tweet may contain up to four). Descriptions may not be added to videos or animated GIFs, the other common media formats on the site, and descriptions may not be added or edited after the tweet is posted. Image descriptions can be added to tweets posted through the API using the “ext_alt_text” tag [31], so third party tools may enable image descriptions for their users as well.

3.2 Snapshot of Image Description Usage

We first aimed to measure the amount of image descriptions on Twitter. Using Twitter’s public API [32] we collected a sample of public tweets across 5 days in June 2018 for an average duration of 12 hours per day. This resulted in a collection of over 9 million tweets from all languages, including both original tweets and retweets (which may be duplicated). Approximately 1.09 million, or 11.84%, of those were tweets with at least one image. Only 1,144, or

0.1%, of those tweets (with images) contained image descriptions for at least one image.

3.2.1 Original Photo Tweets. Upon further examination of the tweets with images, we noted that 271,330, or about 25%, were original tweets. The other 75% were retweets. Only 177, or 0.07%, of original photo tweets contained image descriptions. When we examined the image descriptions in our sample we noticed that some images had URLs as alternative text (often the source of the image), which we filtered out as they are not descriptive. After removing URLs, there were 166 tweets remaining; 0.06% of the original 271,330 photo tweets.

3.2.2 Retweets. If a tweet contains an image description, retweets of it will still contain the description, but new descriptions cannot be added by other users. Retweeted tweets accounted for 75% of the image tweets in our sample, although many of these were duplicated. We had a total of 820,469 retweets, but only 426,084 were unique (51.2%). For unique retweets, 0.11% contained descriptions, which after removing URLs left only 0.05% or 207 tweets.

3.2.3 Photos vs Tweets. Twitter allows a user to attach up to four images to a single tweet, each with its own description. When examining each photo from original tweets we observed that 336,584 photos were shared, and similar to the other categories only 0.05% contained image descriptions (only 0.03% after filtering out URLs).

The details of our sample and breakdown of our analysis is shown in Table 1. Overall, this suggests that less than 0.05% of the image content on Twitter is accessible to screen reader users.

3.3 Accessibility of Popular Accounts

We wondered if, despite the fact that tweets from randomly-sampled users were mostly inaccessible, perhaps more popular accounts would enable and use the feature regularly. These accounts, run by celebrities or organizations, were more likely to have professionals writing their content. Additionally, if popular accounts were accessible, they would have a greater impact on the overall accessibility of Twitter, as they would appear in more user’s timelines.

We collected the image tweets from the top 50 most popular Twitter accounts [33] by number of followers. The tweets collected included every image tweet available between the time the feature was launched and October 2018. We found that only 3 of the 50




Figure 2: The user interface for adding image descriptions on Twitter. The left image shows the “Compose Tweet” window with the “Add Description” pop up circled in blue. The right side is the subsequent window showing the added image with a field at the bottom to add a description.

Table 1: Number of tweets and retweets containing photos and alt-text.

	Day 1	Day 2	Day 3	Day 4	Day 5	Total
Tweets and Photo Tweets						
Total Tweets	1,860,947	1,757,012	1,366,244	2,243,235	1,997,713	9,225,151
Photo Tweets	225,679	189,445	145,778	269,098	261,789	1,091,799
Photo Tweets with alt text	195	199	205	297	248	1,144
Original Photo Tweets						
Original Tweets with photos	57,530	46,793	37,949	65,546	63,512	271,330
Original Photo Tweets with Alt	40	34	24	53	26	177
Photo Retweets						
Photo Retweets	168,149	142,652	107,839	203,552	198,277	820,469
Photo Retweets with alt	155	165	181	244	222	967
Original Photos and Retweeted Photos						
Total Original Photos	71,065	56,736	46,626	81,262	80,895	336,584
Total Retweeted Photos	287,909	241,828	183,244	355,346	349,894	1,418,221

accounts had ever used the feature: the official Twitter account, Bill Gates, and BBC Breaking News. Assuming they enabled the feature just before their first tweet with an image description, these three accounts collectively added descriptions to 14 (8.8%) of 159 image tweets. This assumption does not make much sense for the Twitter account, of course, as it introduced the feature [27]. Only 5 of Twitter’s 44 image tweets since the introduction of the feature have included image descriptions. The other 47 popular accounts that never added descriptions for accessibility to their images included prominent news organizations (New York Times, CNN, ESPN), politicians (Barack Obama, Donald Trump, Narendra Modi), and celebrities (Katy Perry, Justin Bieber, Taylor Swift).

Based on prior work that found government websites more accessible [5, 16], we wanted to see if this trend carried across to government accounts on Twitter. Using a list (created by C-SPAN) of 577 accounts associated with members of the current U.S. Congress [8], we performed the same analysis as above. In all tweets since the introduction of the feature, 42 accounts had used the feature at least once, with an average of 3.8% of their image tweets containing a description.

4 ACCESSIBILITY FOR BLIND USERS

Understanding the state of accessible images on Twitter as a whole is valuable to inform how easy it is for a screen reader user to interact with any piece of content on the platform. However, this is not how most users experience Twitter content. They follow a set of users and only see the content authored, retweeted, or liked by the users they follow.

We analyzed the timelines of 94 self-reported blind users to determine if the experience of using Twitter is more accessible for people with vision impairments. To find blind users (who likely use screen readers to access Twitter), we collected a random sample of the accounts that follow the Twitter account of the National Federation of the Blind (@NFB_Voice), a large US-based organization led by blind people. From this sample, we selected an initial 100 users who self-identified as blind or visually impaired in their Twitter profile description. There is no definitive indication that these users also use a screen reader, but we assume that many do.

We wanted to understand if the timelines of these users had the same level of accessibility as Twitter as a whole, or if the accounts they follow posted accessible content more frequently. For each user, we collected all of the accounts they follow, known as “friends” on Twitter¹. For 6 users, we were unable to retrieve this data, as the accounts they follow were private. For each friend of the 94 remaining accounts, we collected 200 of their most recent tweets (and retweets), or as many as were available. Using these tweets sorted chronologically, we recreated each user’s Home timeline for one day in October 2018.

Table 2: Summary of the tweets in one day of recreated Home timeline for 94 blind Twitter users.

	Min	Max	Average	Median
Total friends	6	5,001	720.7	371
Total tweets in timeline	2	29,231	3,379.7	1,554.5
Percent photos in timeline	0.0%	39.7%	18.4%	18.3%
Percent alt Text in photos	0.0%	41.4%	4.6%	2.0%

This is not a perfect recreation of a user’s Home timeline, as Twitter does not show every tweet chronologically [29]. Additionally, we could not gather tweets from accounts with higher privacy settings (protected accounts), and some posted tweets may have been deleted by the time we collected them. However, we believe this to be a good approximation of the content these Twitter users would have been exposed to if they logged into Twitter that day.

Overall, we found that the recreated timelines for these users included 18.4% of tweets with photos on average, and 4.6% of photo tweets contained image descriptions. Table 2 contains information about the range of content we observed in these users’ timelines, and a visual depiction is shown in Figure 3. In general, these timelines were an order of magnitude more accessible than Twitter as a whole, indicating that these users may be involved in communities with more awareness of the image description feature. An alternative explanation is that these users chose not follow some accounts that post inaccessible images. Regardless of the explanation, while

¹This does not indicate a mutual relationship, as it does on many other social networks.

these timelines are more accessible than our random sample of Twitter as a whole, they were still largely inaccessible.

5 QUALITY OF IMAGE DESCRIPTIONS

After investigating the prevalence of image descriptions on Twitter and finding low usage of the feature, we were curious about the quality of descriptions that do exist. Do authors write image descriptions in a way that is useful for people with vision impairments? Are the descriptions relevant to the photos?

To answer these questions, we developed a four-point rating scale from "Irrelevant" to "Great description" to assess the quality of image descriptions on Twitter posts (Figure 4). We filtered our sample of original photo tweets with descriptions to those in the English language (based on the "lang" attribute in the Tweet metadata), as we could not effectively assess non-English descriptions. This left 93 tweets from 71 users. To get a larger sample of tweets, we downloaded all tweets from these users with alt text.

5.1 Evaluating Image Descriptions

Prior work by Salisbury *et al.* [25] constructed conversations between crowd workers who were not allowed to view the image in a tweet, and crowd workers who were allowed to and expected to describe the image. From this, they built a set of structured questions to help guide composition of alternative text. We merged these questions and guidance from the Web Content Accessibility Guidelines [10] to develop a rating scale for image descriptions on Twitter. We used a rubric to rate the quality of image descriptions, and examples are shown in Table 4.

5.2 Findings

Two of the researchers used the rubric to redundantly code the quality of 500 photo tweets. We estimated the inter-rater reliability for this set by calculating a weighted Cohen's Kappa of 0.83 for these 500 ratings, which can be seen as strong agreement [20]. One of the researchers then rated an additional 500 photo tweets, resulting in a total of 1000 rated tweets.




Figure 3: Each timeline analyzed, with the y-axis being percentage of total tweets. Dark blue: photos as a percent of tweets; lighter bars: alt text as a percent of tweets.

Using this set of 1000 rated tweets, we found that 62.6% of alt text rated as irrelevant or somewhat relevant. Only 15.8% rated as "great", with most elements of the image described. More details are available in Table 3.

5.3 Frequency of Use

Some accounts included in our sample used alternative text often, while others only had included a description in a small fraction of the images they tweeted. We looked at all of the images each account had posted since they first used the image description feature. On average, accounts included image descriptions for 60.5% (median = 66.67%) of images tweeted. Many of the accounts who included image descriptions 90-100% of the time were automatically posting content as bots, however. After manually removing 19 of these accounts (Section 5.4, we found that humans tended to write image descriptions for 49.8% of their photos (median = 50%).

5.4 Quality of Bots vs. People

From this analysis, we noticed that many of the tweets with image descriptions were coming from the same accounts, and these accounts were either explicitly bots or exhibited bot-like behavior. The accounts we knew were explicitly bots indicated they were automated in their name or profile description. Some bots generated specific kinds of memes, such as taking images from museum

Table 3: Percentage of image descriptions rated using our rubric for different samples.

Rating	Sample 1	Sample 2 (no bots)	Combined
Irrelevant	32.6%	1.1%	16.9%
Somewhat relevant	30.0%	15.9%	23.0%
Good	21.6%	34.1%	27.9%
Great	15.8%	48.9%	32.4%
Sample Size	1000	1000	2000

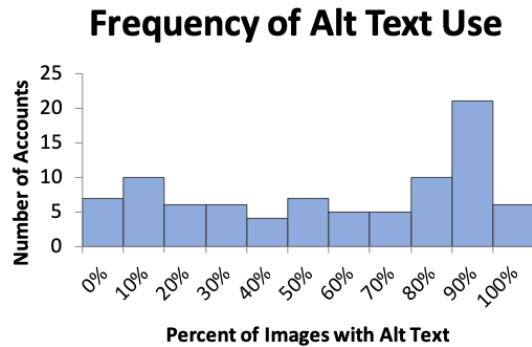



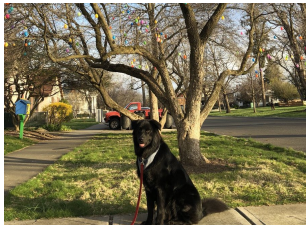


Figure 4: Histogram showing the frequency of image descriptions posted for the accounts in our sample. There is a spike in the 90-100% range that is comprised mostly of bots.

Table 4: Alt-Text Rating Examples. *Links and Usernames are removed from post text for anonymity.

Rating	Image	Post Text & Image Description	Rating and Reason
Irrelevant to image		Post: New Music Video - Beach House (@[username]) "Black Car": [link][link] Alt: Beach House	We rated the alt-text for this image as irrelevant because it does not describe anything in the actual image shown.
Somewhat relevant to image		Post: Trailering training class in Henderson yesterday. Alt: Trailering training class in Henderson yesterday.	We rated the alt-text for this image as somewhat relevant because it relates to the purpose of the image but just repeats the post text without describing more of the image.
Good: some aspects of image described		Post: Collecting all the shirts #DNNSummit Alt: Three t-shirts from DNN summit.	We rated the alt-text for this image as good because it describes most of the things in the image. It could also describe the color and/or what is on each shirt.
Great: almost everything is described		Post: "Taking a moment to appreciate beauty. #DogsOfTwitter @[username] #beauty #spring" Alt: Styx the dog sits in front of a tree with plastic Easter eggs dangling from the branches. It's very pretty.	We rated the alt-text for this image as great because it describes almost everything in the image. This particular description also points out information that might not be immediately apparent visually, either.

archives and superimposing fictional characters on them. Others posted automated information about location-based information such as earthquakes or air-quality, and included maps with their posts. Other accounts did not explicitly state they were bots but exhibited bot-like behavior such as blogs that re-posted articles from their website to drive traffic there from Twitter.

These bots exhibited different levels of accessibility. Some included information that described the entirety of a generated image. Others, such as blog-associated accounts, included human-written alt text for the images in their articles. Overall, however, we saw a general trend of descriptions associated with bot posts receiving a score of "Irrelevant" or "Somewhat relevant".

As we were interested in human-written alternative text, in addition to overall accessibility on Twitter, we sampled an additional random 1000 tweets from our users that excluded the 19 accounts that were explicit or suspected bots. One member of the research team rated this additional sample, finding that just under half of the sample rated as "Great" (48.9%) and 34.1% as "Good". The complete numbers of these samples are available in Table 3.

The difference between these two samples is stark, and indicates that the majority of poor image descriptions may be generated from automated sources. The human tweet authors in our sample described images at a moderately useful level ("good") or higher. Bots are currently important and lively sources of content in the

Twitter ecosystem, and their content should be just as accessible to people with vision impairments. These results for human authors are promising, and show that many users write good image descriptions, if they enable and use the feature.

6 DESCRIPTION AUTHOR INTERVIEWS

We interviewed people who had used the image descriptions feature to add alternative text to images in their tweets. Prior work has demonstrated a clear accessibility need on social media platforms from the perspective of people with disabilities [6, 24, 34]. Twitter users that create content must be able to make their images accessible, as the responsibility of making Twitter and other social media platforms accessible should not fall solely on people with disabilities. We were interested in users' motivations for using the feature, their process for composing image descriptions for tweets, and why they may not always add descriptions to their tweets.

6.1 Participants

We interviewed 20 Twitter users who had written at least one tweet including an image and description. Participants were recruited via direct messages or emails to users identified in our sample from Section 5: Quality Ratings. We also recruited participants via an advertisement on Twitter posted by one of the authors. The mean age for participants was 41.6 years (std. dev.= 11.7) with a

Table 5: Participant Demographics.

PID	Age	Gender	Years On Twitter	Occupation
P1	32	male	10	Make and sell jewelry
P2	32	male	10	Client services
P3	45	male	10	Network Administrator
P4	38	woman	9	Translator
P5	48	female	10	Public Engagement
P6	61	female	10	Dir. of IT Accessibility
P7	22	male	7	Research Fellow
P8	25	male	7	Ph.D student
P9	28	female	4	Research Associate
P10	42	female	9	Photographer
P11	44	female	9	Student
P12	32	male	4	Software Engineer
P13	57	female	9	Retired Hydro-geologist
P14	48	male	6	Construction Manager
P15	45-50	trans	9	Scientist
P16	43	male	12	Software Engineering
P17	55	female	12	Marketing
P18	59	male	8	Engineer
P19	48	male	9	Musician, Developer
P20	32	female	9	Marketing

range of time using Twitter from 4 to 12 years (very early adopters). None of the participants reported any visual impairments or use of screen reader technologies. Thus, participants in our study do not represent visually impaired users on Twitter, but rather represent the creators of content that visually impaired users might consume.

6.2 Interview Format

We conducted interviews in person, over the phone, and through Twitter direct messages. The interview consisted of four topics: activation and use of the image description feature, interactions with blind or screen-reader users regarding alt text, process of writing and examples of their alt text, and suggested changes for the image description feature on Twitter. Demographic information is listed in Table 5. The interview sections are described below.

Feature Use. We asked participants three questions regarding their use of the image descriptions feature: (1) why they activated the feature, (2) how they decide when to add descriptions, and (3) and how often they believe they add descriptions to the images they tweet. Our aim was to understand what motivated these users to activate the feature, how actively they provide image descriptions, and for what purpose.

Interactions with Followers. We asked participants about their interactions with followers who were blind or used screen readers. Specifically, we asked about any direct interactions participants had and how many of their followers they believed benefited from image descriptions.

Examples of Image Descriptions. We discussed participants’ process for composing image descriptions for their tweets. We first asked participants to describe the elements they included in their

descriptions and what they think about when composing them. Next, we chose specific examples from participants’ tweets to discuss; one where they wrote alt text and one where they did not. We discussed how they wrote the specific image description for the example and reasons they may not have added descriptions for other images.

Changes to Image Description Feature. Finally, we asked participants to think of one thing they would change about the Twitter image descriptions feature.

Data and Analysis. Each interview was transcribed and analyzed using a theoretical approach to thematic analysis [7] based on the sections of the interview described above. We coded each transcripts based on the topics of the interview. The first and second authors redundantly coded the first five interviews, then discussed and refined the code book before independently coding the remaining 15 transcripts.

6.3 Findings

6.3.1 Feature Discovery, Motivation, and Use.

Discovery. Participants discovered the image description feature through a variety of ways. The most mentioned (6) was by the suggestion of someone they were following or a tweet they saw mentioning the capability to add descriptions to images on Twitter.

“Because [a specific user I follow] suggested it, I suspect. I can’t say for certain. I would never have looked into the accessibility settings if someone hadn’t said that this feature was hiding there.” – P3

Some participants (2) discovered the image descriptions feature through announcements, presumably made by Twitter, when the feature was released. Three participants anticipated the release of the image description feature before it was announced. P1 had even requested that Twitter add the capability, almost two years in advance of its release.

“I’d been waiting for it to be an option for a while - me and friends had contacted Twitter about it but nothing happened for a while. So as soon as it became available I switched it on...Quite a few of us have disabilities so we try to push for accessibility features like captioning/transcripts or image description.” – P1

Three other participants mentioned they activated the feature because someone else had used alternative text to describe images or they were involved in accessibility related work or communities.

“I learned about it through [work] colleagues and wanted folks using screen readers to be able to access media, too.” –P15

Motivation. In terms of why they chose to use the feature, some participants cited that adding alternative text to images was a low effort activity, but would have a high impact on the people it was intended for.

“I knew I wanted to use it because it’s such a simple small thing that is no effort for me, but may mean a lot to the reader.” – P4

Four participants cited personal or professional connections to someone with a disability as motivation for their use of alternative text. P18 described their professional connections:

“I ran [several] - all accessibility-related accounts. It would have been disingenuous to have the disability community as

our majority base of followers and not include alt-text, so we did as soon as it was available. Same for my personal account since I am in that field and was perhaps more aware of it than others might be." – P18

Two other participants, like P4 and P7, were influenced more personally by their relationship with blind users:

"Maybe subconsciously because a close friend of mine is visually impaired and I used to describe some pictures I find on twitter to him over lunch table conversations but now since we don't meet a lot, this is one way I can keep my friend engaged with my conversations." – P7

Overall, participants demonstrated that, while perhaps not personally connected, providing image descriptions was a matter of "inclusion" that makes things better for everyone.

Habits of Use. With regard to participants' intended use of the feature it became clear that most (12/20) participants explicitly intended to add image descriptions to every image they posted. When we asked how often participants actually used alt text, we found that people experienced varying levels of success:

"Ok I'm not as good at doing alt-text [as] I thought. In my last 10 images, 5 had alt text." – P8

For others, adding the descriptions was a matter of convenience, citing that the applications they used either allowed it or made it straightforward to add the descriptions. Two participants described that when they had time, they would add image descriptions:

"Maybe its really the time which is the factor, when I feel like I am in a rush, I kind of miss adding it to alt-text and instead just describe it in the tweet text." – P7

Two other participants described being reminded by the interface (see figure 2) to add a description.

"The add a description prompt under the photo [which appears once users have activated the ability to write alt text in their accessibility settings] reminds me. I think I do it for every photo I add." – P11

We had two participants, P2 and P19, who created bots that generated original tweets with images that included descriptions. Both bots generated meme-like images by imposing content from online image archives with other images. The alt-text in both cases was just the description of the image that was in the library or museum archive. For these bots, including alternative text as part of the tweet metadata was supported by the Twitter API, thus making the inclusion of a description for every image very convenient.

One participant highlighted that they try to add image descriptions whenever the image is their own photo, whether it is a photo they took or a graphic they created. One participant, P20, mentioned that she switched to using the Twitter web interface or mobile app specifically because another third party tool she was using did not support adding image descriptions. This was similar to P6's professional commitment:

"If I include an image, I add alt text. If I don't feel like bothering, I don't include a picture. Professional pride" – P6

6.3.2 Follower Knowledge and Interaction. Half of participants (10/20) had no knowledge of any of their followers being blind or using screen readers. A majority of participants (11/20) had

never been contacted by a blind person regarding their use of alt text.

"I have put stuff in my alt text to try to solicit a response and never gotten one, so I am not sure any of my twitter followers use screen readers." – P15

The other half of participants (10/20) were aware of at least one blind individual among their followers. A few (8) estimated specific numbers between 1 and 40, or "a small percentage". Participants had limited Interactions with their networks regarding alt-text; however, one participant mentioned having discussions where alternative text was mentioned.

6.3.3 Authoring Image Descriptions.

Description Process. We asked participants to share their process for writing image descriptions, both in general and in relation to a specific example of one of their own tweets. Some participants had specific strategies for describing an image. Depending on the intent of the post (mentioned explicitly by 6 participants), a majority of participants (11/20) described writing a general description of the image. One participant imagines trying to explain an image to a friend on the phone.

One participant mentioned describing the colors that appeared in the image. Others described determining the importance of objects, background, and other content in the image that the reader may not be able to ascertain from the main text of the tweet. Participants mentioned transcribing the text included in images or describing the objects, actions, and facial expressions in the image. Based on the content in the image, three participants tried to highlight the focus of the image.

For bot creators and specific content-focused accounts, participants mentioned that they used the image descriptions to convey the purpose of the tweet. For instance, to describe the important visual elements in the image to represent a joke or meme, or to convey why the image makes the joke funny in the context of the post (examples in Figure 5). We encountered two notable examples of content creators that had difficulty writing descriptions, a photographer and a bot creator with accounts that posts memes. P10 grappled with trying to describe the important photographic elements of her images:

"When it's a photograph being shown as a piece of art, that's where it gets difficult - especially since many of my photographs are quite abstract and tend to defy description! I try to touch on the straightforward visual facts of it (what is it a photograph of) but also get across the sense & feel of it where I can. The latter is in some ways more important with my photography. Things like colour tones (is it cool or warm, soft colours or vivid colours), are there any textures, what does it resemble." – P10

P2, who creates Twitter bot accounts, described nuances of conveying jokes presented through memes:

"[The first bot] deviates from my personal answers since its alt-text is the punchline and doesn't describe the visual content of the meme. That can be hard. For [these tweets] the verbosity is the joke, so I used [the alt-text] to convey that verbosity. For [A different bot account] the joke is both that the character looks intelligent and is making a foolish mistake...the alt text template there mentions the book, the glasses, and then names the two disparate concepts [in the image]." – P2



Figure 5: The alternative text for (a) was: "Cartoon man in glasses holding book looking at a butterfly labeled 'microbes', asking "Is this ponies" For (b) it was: "Monochrome photograph of a number of dandelion seeds tangled together - much of the focus is soft, the the seed heads and fluff are clearly visible in places. There are hints of blue colour tone in the background."

All of these examples illustrate a very nuanced process. The approaches vary from person to person and with the intent of the post and image.

Non-Use. Participants cited time constraints as the most common reason they might have missed adding alternative text, either because it was too time consuming, they were in a hurry and forgot (8), or they had to write multiple tweets in succession and missed adding alt-text in the process (5).

Other reasons were that participants simply forgot what the feature was for or that the feature was there. We encountered four participants who mentioned that they relied on the Twitter interface to remind them to add the description.

"Actually, looking at twitter on web, the cue for alt-text is "description" which sounds pretty optional at a glance. there's a chance i read "description" and forgot it was for the visually impaired in my rush to tweet and accumulate all the engagement.
– P8

Some participants stated that adding image descriptions from a mobile device was still not possible. However, it is currently possible to add descriptions from the Twitter mobile application. This suggests that either participants did not notice the feature or have not used the application since the capability was added.

6.3.4 Changes. The most common suggestion, mentioned by eight participants, to improve the image description feature on Twitter was to make the descriptions visible to sighted users, especially for their own tweets. P4 states:

"[I want] to be able to add the alt text after the tweet is posted. I don't need to edit tweets as a whole, but I would really want to be able to go through all my pics and add alt text. And I wish hovering over images would show what alt text I wrote, like they used to with other images online. Maybe more people would notice/become aware then?"–P4

The only current process for viewing these descriptions requires the user to view the source code for the web page. This suggestion is also in line with the desired ability to edit the tweet content (and image description) after it has been posted, which is typically not supported on the platform.

The second most common suggestion (by 5 participants) was to make the image description feature "active" by default instead of as a setting that you have to turn on.

"I would make it automatically enabled for people so that users don't have to wade through the settings to turn it on. They could be helping so many more people if only they used this feature up front."–P20

Other suggested changes included improving the interface reminder, addressing bugs, and providing automated support to generate image descriptions. Participant 1 also mentioned the need to increase and "normalize" the use of the feature in a similar manner to captioning. In addition to support from Twitter, one participant (P11) wanted access to individuals with vision impairments to provide guidance on writing good image descriptions. Finally, three participants mentioned increasing the character limit (currently 420 characters) for the alternative text to allow more thorough descriptions, especially for screenshots or pictures of text.

7 DISCUSSION

Our analysis of image description on Twitter revealed that very few people (including popular and government accounts) use the feature, with less than 0.1% of original image tweets having any descriptions at all. The image tweets exposed to blind users contain descriptions slightly more often (4.6%), but still are very inaccessible. The Twitter users that do use the feature author descriptions for about half of the images they tweet, and the descriptions they write tend to be "good" or "great" 83% of the time (bots excluded). Our analysis of interviews with image description writers examined the reasons for use (and non-use) of the feature, and lead us to two paths for improvements: those that Twitter (or similar social platforms) could currently undertake, and those that require further research and additional tools.

7.1 Improvements for the Image Description Feature

Interview participants identified many issues with the image description feature that, if fixed, would lead to higher or better usage. Participants requested that alternative text be visible to them and editable after posting. Interview participants had trouble recalling which images they added descriptions to, and what they wrote. The most common reason they did not add a description was that they forgot when posting quickly, and being able to add image descriptions after the fact would be valuable. Taken further, if users could add image descriptions to retweeted images, volunteers or friends of screen reader users could then make this content accessible.

Another common request was for Twitter to just turn on the ability to author image descriptions for everyone, rather than requiring users to find and enable this capability.

"Top wish: it should be turned on by default. It's almost like they're trying to hide it." –P6

Most participants turned on the feature as soon as they found out about it, and try to include image descriptions for every tweet with images that they post. We agree that the image description feature is hard to find and understand, and Twitter should enable it for everyone to increase accessibility on its site.

However, just enabling the feature for everyone is likely not enough. It is currently designed for people who are familiar with image descriptions for people with vision impairments, and instantly enabling it for everyone could lead to misuse. It may be abused to make tweets appear higher in Twitter or external searches or include spam URLs. We have already observed users including subtle messaging in Twitter image descriptions for followers who know to look for it. Even without intentional abuse, new users who do not understand the purpose of the feature or how to write alt text may not produce high quality image descriptions. Twitter should provide clear on-boarding instructions for users when they first use the feature, explaining its purpose and why users should provide image descriptions for their images. In order to reduce confusion when posting a new image tweet, on image upload Twitter should provide instructions on how to write alt text or a template of structured questions, which prior work has found results in higher quality alt text in other media (e.g., STEM textbook diagrams) [23].

7.2 Additional Tooling and Training for Users

We see two major opportunities for researchers to make social media platforms more accessible through additional tools.

Some content on Twitter is ripe for auto-generation of image descriptions. Automatic captions for generic images has been exemplified by ALT text bot [11], which provides automated alt text in response to tweets containing images. However, researchers could go further in areas where the content format is more constrained. Some photos are just photos of tweets, and could be accessible if linked to the original tweet. Screenshots or photos of text are popular (9.7% of the no-bot sample in Section 5), and robust optical character recognition could make these accessible.

Researchers should also develop tools to help users write better image descriptions. Many users do not know what elements to include in an image, and would benefit from specific instructions, such as the structured questions developed in Salisbury *et al* [25]. Automated tooling could rate how descriptive alt-text is, provide specific instructions based on recognized objects in an image, or even pre-fill the image description with an auto-generated scene description [25, 37]. This may help prompt the user to change or refine the image description before publishing.

7.3 Supporting Authorship through Automated Feedback

To provide a starting point for the goal of creating tools for image description authors, we used the sample rated in Section 5 to create an automatic rater for image descriptions. We merged ratings of “Irrelevant” and “Somewhat relevant” alt text as “low-quality” and ratings of “Good” or “Great” as “high-quality”. An Extra Trees classifier [13] was trained on features extracted from a subset of the sample (1,320 tweets). The specific features used were: counts of parts of speech in the alt text and post text, shared words between the alt text and post text, as well as the length of the alt text. This classifier achieved an overall accuracy on the remaining 680 tweets of 85.3% (AUC = 0.84, precision = 0.83, recall = 0.94), demonstrating it is able to distinguish between much of the alt text quality. The five most important features for this classifier were: number of prepositions in alt text, number of words shared between

post text and alt text, length of alt text, number of present verbs in alt text, and number of plural nouns in alt text.

At a very simple level, a classifier like this demonstrates that automatic feedback could be given to users, when they compose their descriptions, on whether it appears to be low or high quality. Specific feedback could focus on how similar the alt text is to the post text, objects in the image that are not mentioned, or a lack of actions (verbs) and objects (nouns) in the written description.

7.4 Limitations

The primary limitation of this work is that we have not comprehensively studied the accessibility of Twitter as a whole, only photo tweets. Other forms of media, such as animated GIFs, videos, polls, and URL previews exist. URL previews in particular can contain alternative text that is pulled from the linked page, but this was not the subject of our analysis.

Our rating scale was developed from prior research on the experience of blind users interacting with alt text, but it was still designed and executed by sighted researchers rating image tweets. Therefore, it does not fully reflect what screen reader users seek when browsing images on Twitter.

8 CONCLUSIONS

Our findings show that image descriptions are very rarely provided on Twitter. This is in large part due to very few users having the feature turned on, but even the users who enabled the ability to provide alternative text descriptions did not always write them. Twitter would likely dramatically improve its overall accessibility if it encouraged all users to provide descriptions and enabled them to do so across all types of media. Access to social media is increasingly important for participation in many aspects of society, including social connection, entertainment, civic participation, and news consumption. By improving descriptions of visual content on social media networks, many users with vision impairments [34] will again have equal access to these vital platforms. As P15 states:

“Only built-in accessibility from the start provides more equitable access. Only doing it because/when someone asks for it or because we know a specific individual needs it, puts the burden on the person with the need, and that’s not how accessibility should work. And in terms of accessibility online, alt text for static imagery is a low hanging fruit, easy and inexpensive to implement. Failure to do it is just inconsiderate laziness.”—P15

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. (DGE 1745016). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

In addition, this work was developed under a grant from the National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR grant number 90DP0061). NIDILRR is a Center within the Administration for Community Living (ACL), Department of Health and Human Services (HHS).

REFERENCES

- [1] Adobe. 2018. Add alternative text. <https://www.adobe.com/accessibility/products/acrobat/pdf-repair-add-alternative-text.html>
- [2] Apple. 2018. Supporting VoiceOver in Your App. https://developer.apple.com/documentation/uikit/accessibility/supporting_voiceover_in_your_app
- [3] Julian Ausserhofer and Axel Maireder. 2013. National politics on Twitter: Structures and topics of a networked public sphere. *Information, Communication & Society* 16, 3 (2013), 291–314.
- [4] Tim Berners-Lee and Dan Connolly. 1995. HTML 2.0 Specification. *W3C: http://www.w3.org/MarkUp/html-spec* 34 (1995).
- [5] Jeffrey P Bigham, Ryan S Kaminsky, Richard E Ladner, Oscar M Danielsson, and Gordon L Hempton. 2006. WebInSight: Making Web Images Accessible. In *Proceedings of the 8th international ACM SIGACCESS conference on Computers and accessibility - Assets '06*. 181. <https://doi.org/10.1145/1168987.1169018>
- [6] Erin L. Brady, Yu Zhong, Meredith Ringel Morris, and Jeffrey P. Bigham. 2013. Investigating the appropriateness of social network question asking as a resource for blind users. In *Proceedings of the 2013 conference on Computer supported cooperative work - CSCW '13*. ACM Press, New York, New York, USA, 1225. <https://doi.org/10.1145/2441776.2441915>
- [7] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (jan 2006), 77–101. <https://doi.org/10.1191/1478088706qp063oa> arXiv:arXiv:1011.1669v3
- [8] C-SPAN. 2018. @cspan/Members of Congress on Twitter. <https://twitter.com/cspan/lists/members-of-congress> [Online; accessed 2-November-2018].
- [9] W3 Consortium et al. 1998. *HTML 4.0 specification*. Technical Report. Technical report, W3 Consortium, 1998. <http://www.w3.org/TR/REC-html40>.
- [10] W3 Consortium. [n. d.]. Web Content Accessibility Guidelines (WCAG) 2.1. <https://www.w3.org/TR/WCAG21/>
- [11] Cameron Cundiff. 2015. alt-text-bot: automatic text descriptions of images on Twitter. <https://devpost.com/software/alt-text-bot>
- [12] Daniel Dardailler. 1997. The ALT-server ("An eye for an alt").
- [13] Pierre Geurts, Damien Ernst, and Louis Wehenkel. 2006. Extremely randomized trees. *Machine learning* 63, 1 (2006), 3–42.
- [14] Google. 2018. Content labels. <https://support.google.com/accessibility/android/answer/7158690?hl=en>
- [15] Darren Guinness, Edward Cutrell, and Meredith Ringel Morris. 2018. Caption Crawler: Enabling Reusable Alternative Text Descriptions Using Reverse Image Search. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 518, 11 pages. <https://doi.org/10.1145/3173574.3174092>
- [16] Stephanie Hackett, Bambang Parmanto, and Xiaoming Zeng. 2004. Accessibility of Internet Websites Through Time. In *Proceedings of the 6th International ACM SIGACCESS Conference on Computers and Accessibility (Assets '04)*. ACM, New York, NY, USA, 32–39. <https://doi.org/10.1145/1028630.1028638>
- [17] Xiaodong He and Li Deng. 2017. Deep Learning for Image-to-Text Generation: A Technical Overview. *IEEE Signal Processing Magazine* 34, 6 (nov 2017), 109–116. <https://doi.org/10.1109/MSP.2017.2741510>
- [18] Instagram. 2018. Creating a More Accessible Instagram. <https://instagrampress.com/blog/2018/11/28/creating-a-more-accessible-instagram/>
- [19] @JohnMu. 2018. Alt text is extremely helpful for Google Images – if you want your images to rank there. Even if you use lazy-loading, you know which image will be loaded, so get that information in there as early as possible & test what it renders as. <https://twitter.com/JohnMu/status/1036901608880254976>
- [20] Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica: Biochemia medica* 22, 3 (2012), 276–282.
- [21] Microsoft. 2017. Seeing AI | Talking camera app for those with a visual impairment. <https://www.microsoft.com/en-us/seeing-ai/>
- [22] Microsoft. 2018. Add alternative text to a shape, picture, chart, SmartArt graphic, or other object. <https://support.office.com/en-us/article/add-alternative-text-to-a-shape-picture-chart-smartart-graphic-or-other-object-44989b2a-903c-4d9a-b742-6a75b451c669>
- [23] Valerie S. Morash, Yue-Ting Siu, Joshua A. Miele, Lucia Hasty, and Steven Landau. 2015. Guiding Novice Web Workers in Making Image Descriptions Using Templates. *ACM Trans. Access. Comput.* 7, 4, Article 12 (Nov. 2015), 21 pages. <https://doi.org/10.1145/2764916>
- [24] Meredith Ringel Morris, Annuska Zolyomi, Catherine Yao, Sina Bahram, Jeffrey P. Bigham, and Shaun K. Kane. 2016. "With most of it being pictures now, I rarely use it". In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. 5506–5516. <https://doi.org/10.1145/2858036.2858116>
- [25] Elliot Salisbury, Ece Kamar, and Meredith Ringel Morris. 2017. Toward Scalable Social Alt Text: Conversational Crowdsourcing as a Tool for Refining Vision-to-Language Technology for the Blind. *Aaai Hcomp 17 Hcomp* (2017), 147–156. www.aaai.org https://www.microsoft.com/en-us/research/wp-content/uploads/2017/08/scalable_social_alttext.pdf
- [26] Hironobu Takagi, Shinya Kawanaka, Masatomo Kobayashi, Takashi Itoh, and Chieko Asakawa. 2008. Social Accessibility: Achieving Accessibility Through Collaborative Metadata Authoring. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility (Assets '08)*. ACM, New York, NY, USA, 193–200. <https://doi.org/10.1145/1414471.1414507>
- [27] @todd. 2016. Accessible images for everyone. https://blog.twitter.com/official/en_us/a/2016/accessible-images-for-everyone.html
- [28] Andranik Tumasjan, Timm Oliver Sprenger, Philipp G Sandner, and Isabell M Welpe. 2010. Predicting elections with twitter: What 140 characters reveal about political sentiment. *Icwsn* 10, 1 (2010), 178–185.
- [29] Twitter. 2018. About your Twitter timeline. <https://help.twitter.com/en/using-twitter/twitter-timeline>
- [30] Twitter. 2018. How to make images accessible for people. <https://help.twitter.com/en/using-twitter/picture-descriptions>
- [31] Twitter. 2018. Post, retrieve and engage with Tweets. <https://developer.twitter.com/en/docs/tweets/post-and-engage/api-reference/get-statuses-lookup.html>
- [32] Twitter. 2018. Sample realtime Tweets. <https://developer.twitter.com/en/docs/tweets/sample-realtime/guides/recovery-and-redundancy>
- [33] Wikipedia contributors. 2018. List of most-followed Twitter accounts – Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=List_of_most-followed_Twitter_accounts&oldid=866718146 [Online; accessed 2-November-2018].
- [34] Shaomei Wu and Lada A. Adamic. 2014. Visually impaired users on an online social network. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*. ACM Press, New York, New York, USA, 3133–3142. <https://doi.org/10.1145/2556288.2557415>
- [35] Shaomei Wu, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. Who Says What to Whom on Twitter. In *Proceedings of the 20th International Conference on World Wide Web (WWW '11)*. ACM, New York, NY, USA, 705–714. <https://doi.org/10.1145/1963405.1963504>
- [36] Shaomei Wu, Jeffrey Wieland, Omid Farivar, and Julie Schiller. 2017. Automatic alt-text: Computer-generated image descriptions for blind users on a social network service. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 1180–1192.
- [37] Yuhang Zhao, Shaomei Wu, Lindsay Reynolds, and Shiri Azenkot. 2017. The Effect of Computer-Generated Descriptions on Photo-Sharing Experiences of People with Visual Impairments. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 121 (Dec. 2017), 22 pages. <https://doi.org/10.1145/3134756>