The Effects of “Not Knowing What You Don’t Know” on Web Accessibility for Blind Web Users

Jeffrey P. Bigham
HCI Institute
Carnegie Mellon University
jbigham@cs.cmu.edu

Irene Lin
HCI Institute
Carnegie Mellon University
iwl@andrew.cmu.edu

Saiph Savage
HCI Lab
West Virginia University
saiph.savage@mail.wvu.edu

ABSTRACT

Web accessibility and usability have been extensively studied for blind web users. The focus has generally been on making it technically possible for blind users to access content, or on helping to make the web more usable. This paper explores a challenge at the intersection of these two lenses, which is the effects of blind web users not knowing what they don’t know. On the web, this often means that the user is having a problem completing a task, but does not know whether the problem is because the information is there and not accessible, whether the information is simply difficult to access, or whether the information is not present at all. We first discuss how this issue has manifested itself in other work in this space. We then present the results of a study with 30 sighted web users and 30 blind web users exploring the phenomenon, demonstrating that not knowing the source of a problem causes frustration and wastes time. We conclude with recommendations for future research to help understand and address this problem, as well as design implications for future technology that may assist non-visual web navigation.

CCS Concepts

•**Human-centered computing** → *Accessibility theory, concepts and paradigms; Empirical studies in accessibility;*

Keywords

accessibility; web; usability; blind; screen reader

1. INTRODUCTION

Accessing the web is a vital part of modern life, and remains more difficult for blind web users [18]. The focus on this problem has been on making the web more accessible, with guidelines [15, 24], assessment tools and processes [29, 26, 19], and automatic and crowd-powered fixing [27, 8]. These efforts have had mixed success [17]. Yet, we know that many of the barriers that blind web users experience are not due to accessibility problems, *per se*, but rather usability problems that have the effect of making content much more difficult to access (effectively inaccessible) [22].

ASSETS ’17, October 29-November 1, 2017, Baltimore, MD, USA

ISBN 978-1-4503-4926-0/17/10. . . $15.00

DOI: https://doi.org/10.1145/3132525.3132533

One such problem that blind web users encounter while using the web is that it is often unclear to them what information is available, and what is available but inaccessible [24]. An example is a phone number on a restaurant’s web site that is in an image without an alternative text description. The information is on the web site, but as a blind user without access to the image, you wouldn’t be able to know for sure (let alone get the phone number). Information that you would expect to find on a web page is often not there[1]. Some restaurants really don’t have their phone number on their web page.

Knowing whether the information you want is on the page is important because it determines the best action for you to take: if the information is not there, you should look for it elsewhere; whereas, if the information is there but inaccessible, it might make sense to employ more costly strategies for accessing it [9], request a fix [27, 8], or find a sighted person to help [4, 7]. Given the pervasiveness of technically inaccessible content, many blind web users develop such strategies for accessing this content.

We call the generalization of this problem “not knowing what you don’t know" (NKWYDK). We have hypothesized that blind web users have greater difficulty in recognizing when expected information is not present, in part because sometimes that information is present but inaccessible. A consequence could be that while a sighted user can quickly move on once encountering inaccessible information, a blind user takes longer or becomes more frustrated.

In this paper, we explore NKWYDK with 60 participants, 30 sighted and 30 blind, in a controlled study. We created tasks that were designed to be easy, hard, technically inaccessible, and impossible for blind users (described in more detail later). We note that sighted users are faster and more successful across easy and hard tasks, as would be expected from prior work [6]. We also find that sighted users are (not surprisingly) more successful in solving the inaccessible tasks. Some blind participants were particularly resourceful in solving them as well, confirming that “inaccessible” does not always mean impossible. Underscoring the point about NKWYDK, we find that blind participants spent similar amounts of time on inaccessible and impossible tasks, and often incorrectly assigned blame in both cases, i.e., thought inaccessible tasks were impossible, and impossible tasks were inaccessible.

2. RELATED WORK

The traditional way to think about accessibility for the web has been in terms of compliance with accessibility standards. The most popular standard is the World Wide Web Consortium (W3C)’s Web Content Accessibility Guidelines (WCAG). The first version, WCAG 1.0, was defined in 1999 [15], and was updated to WCAG 2.0 in 2008 [13]. While there are substantial differences between the

[1]XKCD humorously illustrated this in the case of university web pages: https://xkcd.com/773/

two [24], both are comprised of a number of guidelines and checkpoints within those guidelines that is each assigned a priority level. WCAG has formed the basis for much of the web accessibility legislation that has been passed in various countries, and, thus, achieving compliance is a significant part of what tools test for and what accessibility professionals help to do [12].

Standards such as WCAG fail to adequately capture many impactful usability problems [22, 24]. This is problematic because severe usability problems have similar effects as inaccessible content on a blind web user: the task intended to be completed cannot be successfully completed. Blind people use various strategies to overcome accessibility and usability problems on web pages [9]. Strategies for making the web more efficient to navigate include using CTRL+F to find relevant content, and browsing through a list of links on a page instead of its full content. Strategies for overcoming accessibility problems include using the screen reader's built-in virtual cursor to explore visual elements on the web page, or even using an image's filename to guess at its contents.

Effectively employing these strategies is difficult because *(i)* they do not always work, and *(ii)* using them effectively generally requires some technical knowledge about how web pages are built and how the screen reader interprets them. If NKWYDK causes greater uncertainty about what is on a web page, it may make it more difficult to employ such strategies effectively. Viewed through the theoretical lens of information foraging [23, 14], blind web users likely miss out on useful information because of NKWYDK as compared to sighted users.

Power *et al.* studied the problems that blind web users encountered on the web that were not captured by standards like WCAG [24]. Major categories of problems not covered by WCAG included: *(i)* content found in pages where not expected by users, *(ii)* content not found in pages where expected by users, *(iii)* pages too slow to load, *(iv)* no alternative to document format (*e.g.*, PDF), *(v)* information architecture too complex (*e.g.*, too many steps to find pages), and *(vi)* broken links. At least *(i)* and *(ii)* could be interpreted through the lens of NKWYDK, and these two categories collectively accounted for 13.5% of all problems that users encountered in their study. In particular, the *content not found in pages where expected by users* category was described as follows:

> *"...describes problems where users confidently followed a link to a page, but a piece of information that they expected to find there was missing. For example, on a museum website, users followed a link to an object in the museum collection, but did not find any information about the room in which that object is displayed, which they expected."*

Given that this is already a known problem, we might expect the problem to be even worse if it turned out blind users could not reliably tell if they had encountered it or not. Blind web users seem to behave in ways that one might expect if they were unsure as to whether desired information was on a page they visited. For instance, a behavior noticed in the WebinSitu study was what the authors called "probing," which was when participants visited a web page and then quickly left [6]. Blind participants exhibited this behavior much more than sighted participants did, potentially because they were using the behavior to resolve uncertainty about what the web page contained. Given that redundant information exists on the web, a useful strategy may be to visit web pages, quickly determine if they satisfy your information need, and move along to the next web page if they do not. It may make sense to employ even imperfect heuristics if the cost of trying to find non-existent information is high enough (as seems likely).

NKWYDK also has implications for technology designed to fix accessibility problems based on user requests. For instance, the crowdsourced Social Accessibility project could fix web accessibility problems for blind web users, but the crowd would only be asked to fix problems when requested by users. They note that many of their users had difficulty knowing when and what to ask due to uncertainy about whether problems encountered were because of accessibility or not [28]:

> *"One of the most important findings of the service was the high ratio of users who are not aware of the nature of the problems they are facing. Skilled and experienced users are aware of their problems, and they can easily report them. However, we were surprised by how many users were unaware of the real causes of their difficulties."*

Similar observations were made by Petrie *et al.* in noting that remote participants don't always know why they are experiencing problems while running remote users studies with blind participants [21]. We might reasonably expect that blind users therefore under-report accessibility problems to the owners of web sites, even when companies have methods for reporting problems [4, 10].

The phenomenon of NKWYDK is not restricted to blind web users. NKWYDK naturally comes up in the context of accessibility, especially for those with perceptual disabilities. For instance, one of the use cases described in Scribe4Me was for deaf people to use the system to retroactively understand what hearing people were reacting to [20]. In effect, it was a tool for helping deaf people query into the past for what they didn't know. However, they would only know to make such a query when something in their environment signaled that they should ask. One of the primary challenges that the visual question answering service VizWiz ran across was blind users not knowing whether they had captured the necessary information to answer their question in the photos that they took [11]. NKWYDK also causes problems for people without disabilities, often when making decisions for which some information is unknown (because of time or other costs in retrieving it) [25].

3. METHODOLOGY

To explore the effects of not knowing what you don't know on web usability, we designed a study, which was conducted remotely. Prior research has shown that remote studies can be effective for people with disabilities [21]. Some problems will go under-reported, and so our study was designed to explicitly consider this problem – we wanted users to deal with the uncertainty about the problems they faced independently. Remote studies may be more ecologically valid for blind web users because each participant uses their own screen reader and settings [6].

The study asked participants to complete 8 information finding tasks. It included 4 difficulty conditions (Easy, Hard, Inaccessible, and Impossible). Participants completed 2 tasks in each condition across a total of 8 trials. Task difficulty was counterbalanced and presented in a random order. Each participant completed only one task on each of 8 different web pages to avoid learning effects, with the assignment of condition to web page also randomized. Participants were informed of what the study entailed and signed a consent form prior to participating in the study.

3.1 Web Pages and Tasks

Eight web pages were selected for the study (Figure 1). These web pages were chosen so that participants were unlikely to have visited them before (and thus were unlikely to have been familiar

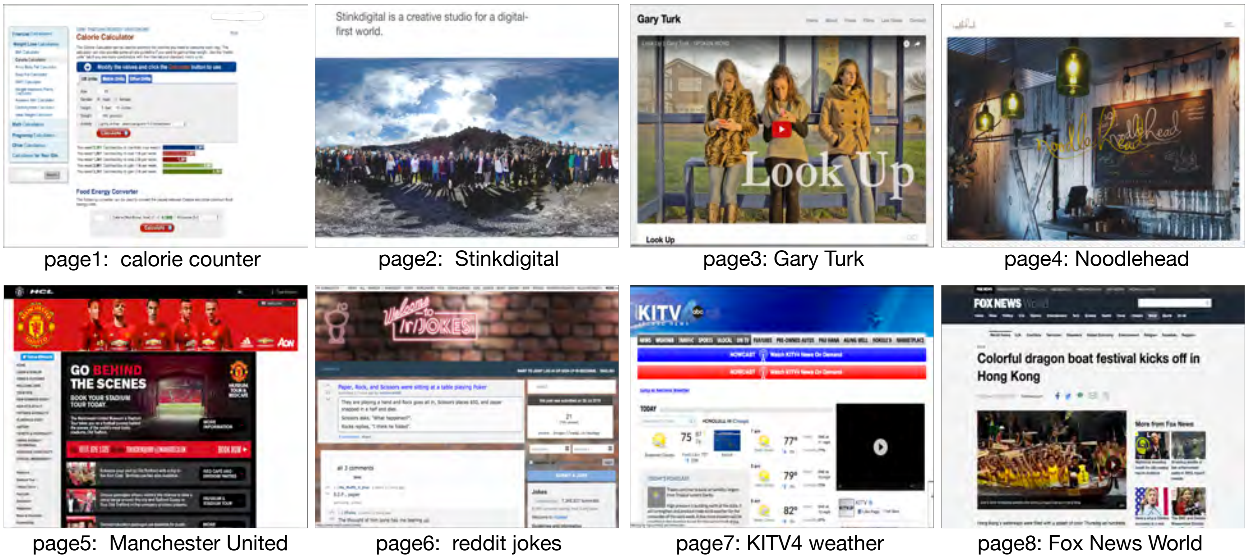

Figure 1: The tasks that participants were asked to complete were created using saved versions of the eight web pages above. All were real web pages, chosen so that participants were unlikely to have experience using them. They were slightly modified as necessary to create tasks of the four types (Easy, Hard, Inaccessible, and Imposible). The full task pages are included with this paper as supplementary material and are described in detail in Table 1.

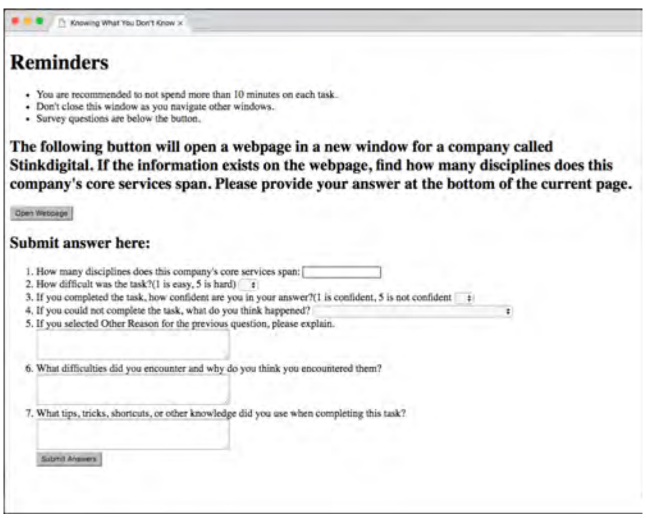

Figure 2: An example task page from our study, which asks participants to answer a question on the given page, times how long they take to answer, and asks them questions about the task.

with their structure)[2]. The web pages chosen varied in accessibility according to WCAG 2.0, with the most common problems relating to lack of alternative text for each image, lack of proper use of headings, and lack of form labels. For each page, we created four tasks with each of the following difficulties, sometimes modifying the page slightly, but often using the page as it was:

- **Easy:** information can be found without major obstacles; nearly all participants should be successful.
- **Hard:** not easily searchable, hidden, or requires an aggregate view (or skim) of the page; many participants should be successful, but it may require longer.
- **Inaccessible:** requires access to content that is inaccessible to screen reader users according to WCAG; some blind participants may still get these correct, but not many.
- **Impossible:** the requested information doesn't exist on the page; no one should get this correct.

[2]The task pages have been included as supplementary material.

In order to reduce bias when users approached a task, instructions were designed to avoid blind participants from assuming "the information will be there." This might have caused participants to spend more time than they might otherwise. The following points were made on the instructions page in addition to other clarifying instructions:

- Some tasks are difficult or impossible. Please spend no more than ten minutes trying to find the answer.
- It is acceptable to answer "I don't know" if you do not have an answer.

Participants were aware that they would receive no penalty for not being able to find an answer, although all participants typed something into the answer box, whether it was a guess or simply “I don‘t know,” and submitted it. The directions also gave a recommended time allotted for each question of no more than 10 minutes. After this time, an alert would pop up asking the participant to move on to the next task. This was to encourage the participant to move on if they got stuck on a particular task, especially given that some tasks were not possible to complete. Participants were able to terminate their session whenever they chose and were encouraged to take breaks between tasks, as needed. They were, however, asked to complete each individual task without interruption.

3.2 Data Collection

We collected three types of data from participants in our study while they completed the 8 tasks. The first was the answer to the questions that we asked, the second was automatic timing data, and the third was a set of questions designed to elicit what participants thought of each task they completed (Figure 2). Task completion time was measured as the time between when the participant initially opened the webpage to begin searching for an answer (when

Task	Description	Explanation
Easy - information can be found without major obstacles		
page1-ez	How many calories are needed to lose 1 pound per week?	Located at the top of the page.
page2-ez	How many disciplines does this company's core services span?	Can be located with CTRL+F.
page3-ez	Name the person who filmed and edited the video?	Located at the top of the page.
page4-ez	How many days a year is this restaurant open?	Located at the top of the page.
page5-ez	This web page provides a link to what social media?	Located at the top of the page.
page6-ez	How many comments are there?	Mentioned twice.
page7-ez	What is the UV index?	Can be located with CTRL+F.
page8-ez	When was this article published?	Logically located after the article title.
Hard - not easily searchable, hidden, or requires an aggregate view (or skim) of the page		
page1-hrd	How many kilojoules per ounce are there in ethanol?	Located in unmarked table.
page2-hrd	Name one award Stinkdigital earned?	HTML (DOM) is not in a logical order.
page3-hrd	What is the length of the video?	Requires pressing play and knowledge of where information would be located.
page4-hrd	What is the price of beef soup?	No heading tags, only CSS styling.
page5-hrd	Name one prize this business boasts on its webpage?	No obvious label.
page6-hrd	What does the sub-comment say?	No obvious distinction between comment and subcomment.
page7-hrd	What is the low in Honolulu for today?	Styling error causes most screenreaders to overlook.
page8-hrd	What is the image caption for the first image?	No <img> tag; the caption is not obvious and blends into the main article.
Inaccessible - technically inaccessible with a screen reader user because of something that is inaccessible on the page		
page1-in	How many calories are needed to lose 1 pound?	Uses inaccessible HTML <canvas>.
page2-in	Name a client of Stinkdigital?	Client logos are displayed as background images.
page3-in	Name the person who did sound engineering for the video?	Requires mouse click because of mousedown event.
page4-in	Name one $8 noodle dish that is indicated as spicy?	Spiciness is indicated by color styling only.
page5-in	What is the email address to contact regarding stadium tours?	Information is contained in an unlabeled image.
page6-in	How many upvotes does the original post have?	Uses aria-hidden tag.
page7-in	When is the sunrise in Honolulu for today?	Requires mouse click because of mousedown event.
page8-in	Find what the image caption about the duck in the river says?	Located in the last slide of carousel of images.
Impossible - information doesn't exist on the page		
page1-imp	How many calories are needed to lose 1 pound?	Does not exist.
page2-imp	Name a client of Stinkdigital?	Does not exist.
page3-imp	Who did sound engineering for the video?	Does not exist.
page4-imp	Name one $8 noodle dish that is indicated as spicy?	Does not exist.
page5-imp	What is the contact email address for stadium tours?	Does not exist.
page6-imp	How many upvotes does the original post have?	Does not exist.
page7-imp	When is the sunrise in Honolulu for today?	Does not exist.
page8-imp	What does the image caption about the flag say?	Does not exist.

Table 1: The 32 tasks of varying difficulty across 8 web pages along with brief descriptions explaining how each met the requirements of its condition. Tasks and problems were based on those found naturally on the selected pages, although in some cases pages were altered to make content accessible or inaccessible as appropriate for the given difficulty.

the user clicked the "Open Webpage" button in Figure 2) to when the participant began inputting a response in any of the required form fields. Data was stored per user in a text file using PHP. Participants were also asked to complete a pre-survey and a post-survey asking about their experience using the web and their experience completing the study, respectively.

4. STUDY WITH SIGHTED PARTICIPANTS

To validate our apparatus and provide a comparison point to our later study with blind participants, we first ran our study with 30 sighted participants recruited from Amazon Mechanical Turk (Table 2). Participants were required to have completed at least 1000 HITs and have a 95% approval rate on the platform. Workers were paid $5.00 for their participation, which for sighted participants averaged less than 15 minutes, resulting in an average hourly wage of more than $20.00/hour.

One goal of running the study first with sighted participants was to validate our study apparatus. We expected that sighted participants would answer most questions correctly across the first three conditions (Easy, Hard, and Inaccessible) but would not be able to answer the Impossible question. The results from sighted participants agreed with these expectations (Figure 3). Overall, sighted participants answered 88.6% of the Easy, Hard, and Inaccessible questions correctly, and failed to answer any of the the Impossible questions correctly (suggesting the answers could not be easily guessed, and were not accidentally on the page). Sighted participants also took longer to complete the Impossible tasks. It required sighted participants an average of 110.57 seconds (SD = 66.55 s) to

Years of Web Experience	# of Participants
5 years or less	0
Between 6 - 10 years	5
Between 11-15 years	5
Between 16-20 years	17
21+ years	3

Table 2: Sighted User Web Experience Levels

complete Impossible tasks, compared to 84.2 seconds (SD = 105.04 s) to complete Inaccessible tasks, presumably because they spent time trying to find the answer before giving up at some point.

When sighted participants were asked in the pre survey about when they felt confused on the web, most sighted participants said that they do not often get confused. One sighted participant responded, *"Not particularly [confused often]. I believe by enough searching that you can find the answer to anything and everything."* When sighted users could not find information, they tended to try modifying their search terms or believed they were looking in wrong place. One sighted participant wrote, *"If the information is not in the place I would expect it to be I get a bit confused. Or if the information is in a format I do not understand."* It was rare for sighted participants to believe that it was the website's fault too. Most sighted participants indicated a solution similar to this user's approach: *"I try harder to find it."*

5. STUDY WITH BLIND PARTICIPANTS

For this study, we recruited 30 blind participants by advertising our study on Twitter. Participants were paid $35.00 for their participation, which required approximately one hour of work. We regularly pay more for blind participants who have proficient skill in using a screen reader, just as we would pay more for other rare individuals with specialized skills.

In the presurvey,many blind participants mentioned that they had trouble identifying the root of problems which they experienced on the web, for instance, *"Yes I do feel this way because sometimes I am not sure if certain links are hidden or if certain content on the page is difficult to find."* Some commented on specific aspects of a website layout, such as an overabundance of links, misuse of ARIA labels or information not presented in an accurately logical and orderly manner. One participant's response ties these comments about the layout together: *"I have found a growing trend toward information being buried in graphics or only available when clicking, tapping or otherwise moving some other part of the interface. The model of a static web 'page' with everything it has to say being said in a straightforward fashion seems to be changing into a far more dynamic and, often, less usable experience for me."* A few users said that since they consider themselves accessibility experts, they don't get confused because they believe the presence of information is binary: *"it's either there or not."* An overwhelming majority indicated that when choosing a website to navigate, content matters more than accessibility. As one user said, *"If a website has better content and is accessible enough, then I will put up with minor accessibility issues to get access to that superior content."*

5.1 Task Completion Times and Accuracy

Overall, it required blind participants 107.1 seconds (SD = 65.07s) on average to complete Easy tasks, 189.9 seconds (SD = 140.32s) to complete Hard tasks, 172.5 seconds (SD = 105.62s) to complete Inaccessible tasks, and 223.7 seconds (SD = 146.35s) to complete Impossible tasks. We conducted a two-way ANOVA to study the effects that the type of task *i.e.*, Easy, Hard, Inaccessible, or Impossible, and being either a sighted or blind individual, had on the time it took people to complete the task (*i.e.*, the type of task and being blind/sighted were the independent variables). There was no significant difference between the amount of time that a blind/sighted participant took to do the task given the task category (F=1.05, p<.38). In other words, the time differences observed between sighted/blind people on a particular task category appear to be explained by being sighted/blind, rather than by the task category itself.

Years of Screen Reader Experience	# of Participants
<= 5 years	9
6 - 10 years	6
11-15 years	4
16-20 years	8
21+ years	2

Screen Reader Used	# of Participants
JAWS	14
NVDA	12
VoiceOver	5

(Note: Some participants use more than one screen reader)

Table 3: Experience and screen reader preferences of blind participants. In this study, we define "blind web users" as those who primarily use a screen reader to browse the web.

In order to study in greater detail how the task difficulty affected the task completion time for blind participants, we conducted an ANOVA test. This test considered for blind people the effect of time to complete the task given the difficulty, *i.e.*, Easy, Hard, Inaccessible, or Impossible. The Bonferonni correction accounted for repeated tests [1]. A significant difference emerged between the amount of time that a blind participant took to complete the task given the difficulty (F=9.87, p<0.0001). Most of the pair-wise task completion times were also significantly different at the $p < 0.01$ level: Easy vs. Hard, Hard vs. Impossible, Inaccessible vs. Hard, and Inaccessible vs. Impossible.

Nearly all blind participants (95.1%) were able to successfully complete the Easy tasks, whereas only 37.7% of blind participants successfully completed the Hard tasks, and 24.6% of blind participants completed the Inaccessible tasks. The distributions of the Easy and Hard conditions differed significantly (Wilcoxon signed-ranks, $Z = -5.0162, p < 0.01$). The distributions of the Hard and Inaccessible conditions were not detectably significant (Wilcoxon signed-ranks, $Z = -1.3207, p = .18$).

Some users were able to use various strategies to correctly complete 4 of the 8 Inaccessible tasks even though it was technically inaccessible according to the guidelines. 5 answered *page3-in* by managing to click the button to display the dynamic content, and then finding the desired information; 1 answered *page4-in* by using built-in screen reader functions to determine the color of the font; 1 answered *page5-in* by using an OCR plug-in; and 8 answered *page7-in* by using the built-in screen reader search function to find hidden content that happened to list the correct answer, although for a day other than the one asked about. We explore these clever solutions in greater detail in Section 5.2.

Given that some of the technically inaccessible tasks were nevertheless answered correctly by some users, primary analysis was repeated using only the tasks that were not solved by any users (tasks 1,2,6, and 8). The average time for correct answers was 130.20 seconds (SD = 103.76), while the average time for incorrect answers was 188.38 seconds (SD = 102.12). Thus, the inaccessible tasks that users could not solve were closer to the average time required to solve those tasks that were impossible.

Participants generally found the desired information when completing the Easy and Hard tasks. Blind users took an average of

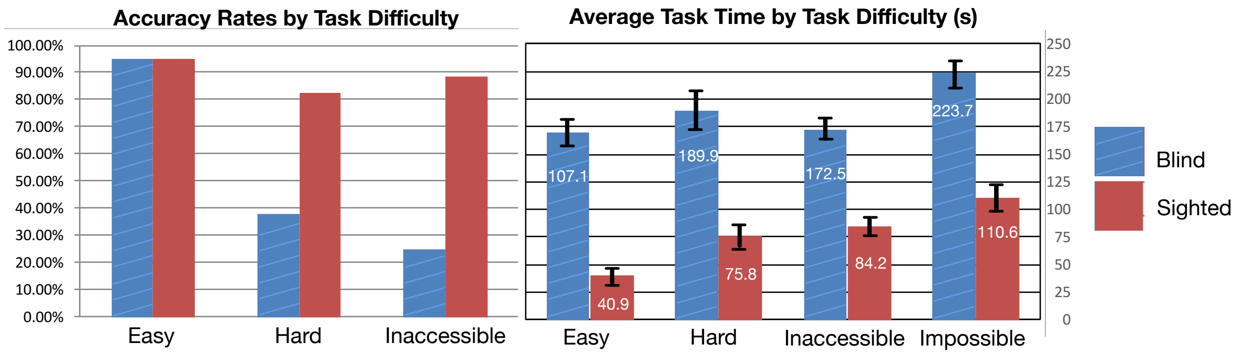

Figure 3: Accuracy and time for blind and sighted participants across the four difficulty conditions. Sighted participants were fairly accurate in all possible conditions, whereas blind participants did well in the Easy tasks, but much worse while adapting the Hard (37.7%) and Inaccessible (24.6%) conditions. Blind participants answered some fraction of the Inaccessible tasks correctly for four of the eight tasks, which we explain in the text. Overall, blind participants required approximately twice as long to complete tasks in each condition in comparison with other participants. Tasks in the Hard condition took slightly longer than tasks in the Inaccessible condition.

107.14 seconds (SD = 65.07s) to complete Easy tasks and average 189.91 seconds (SD = 140.32s) to complete Hard tasks. Blind users had more difficulty isolating the cause of problems in the Inaccessible and Impossible conditions (average of 172.45 seconds (SD = 105.62 s) to complete Inaccessible tasks, compared to 223.7 seconds (SD = 146.35s) to complete Impossible tasks.)

After completing each task, participants were asked to answer the following question: *"If you could not complete the task, what do you think happened?"* Five options were provided: 1) "Not Applicable"; 2) "The information was hidden or not easily searchable"; 3) "I cannot access the information with my screen reader"; 4) "The information does not exist on the page"; and 5) "Other Reason." This question was asked in order to gauge the participant's perceived difficulty level of the task. The first four options respectively represent Easy, Hard, Inaccessible, and Impossible. "Other Reason" provides an alternative to the other options if the participant felt that their response did not fall into any of the first four categories. Sighted participants were able to figure out the cause of problems they were experiencing, but blind participants had much more trouble, often confusing the Hard, Inaccessible, and Impossible categories. Results are summarized in Figure 4.

5.2 Who Receives the Blame

After each task, blind users were asked to explain what they believed was the cause of any experienced difficulties in navigating the web. Researchers used open coding to analyze every blind participant's survey response. Themes emerged from the iterative coding and refinement of the responses.

5.2.1 Bad HTML Design

This category is about blaming the design of the website for the problems experienced. Blind participants thought that faulty HTML was to blame for the problems they experienced 31% of the time. In this instance, the user did not consider that the problem they were experiencing was because they were blind, but rather it was a problem experienced by all users of that site (including sighted individuals.) For example, one participant was not able to find a price on a website, and blamed the page layout:
"*There were too many things to click on, and I must admit I didn't have the patience to check every single thing the prices could have been hidden under.*"

Another participant was unable to find an email on a website and thought it was because most people design websites to not include an email address:
"*...I just didn't find an Email address. The contact link didn't reveal it either. I rarely do find E-mail addresses, so I thought they weren't often provided...*"

5.2.2 Web Accessibility Problems

25% of the time, blind users believed that the obstacles they experienced came from websites that had accessibility issues. In this case, users thought that the designers of the page had not made certain elements available to them, considering only that sighted users might be able to access the information. For instance, this user mentions how he believes the information he is looking for is in an inaccessible image: "*I think the required information was written in a not labelled graphic.*"

A few blind users were able to access the unaccesible information from website through several different hacks, and then stated that it was actually an easy task to do. For instance, the following person was asked to find how spicy was a certain dish, but spiciness of a dish was embedded in a chili icon that has no alternative text. The person hacked a solution that studied the page's html to figure out what was the level of spiciness.

> *"Spicy dishes are indicated using an alternative font size/colour. To figure this out, I had to locate the "Indicates Spicy" text, and check how the formatting of it differed from the surrounding text on the page. I then had to check the formatting for every $8 dish to determine which ones were spicy."*

5.2.3 Self Blame

This category is about blind users blaming themselves for not being able to access or obtain certain information. 13% of the time, people did not blame the website's design or the accessibility of the site, but rather their own abilities. "*I'm sure this information is there but I couldn't find it.*" 31% of the time, blind users did not blame anyone or anything else for issues they experienced.

5.3 Power Users

In the individual task feedback, we received a number of interesting responses from users, who offered clues about how they either did solve some of the tasks in the study, or how they would generally approach these problems. Some responses and behaviors indicated a level of technical expertise necessary to effectively overcome challenges in what seem like routine challenges. Given that these users were among those who seemed to adapt to the prob-

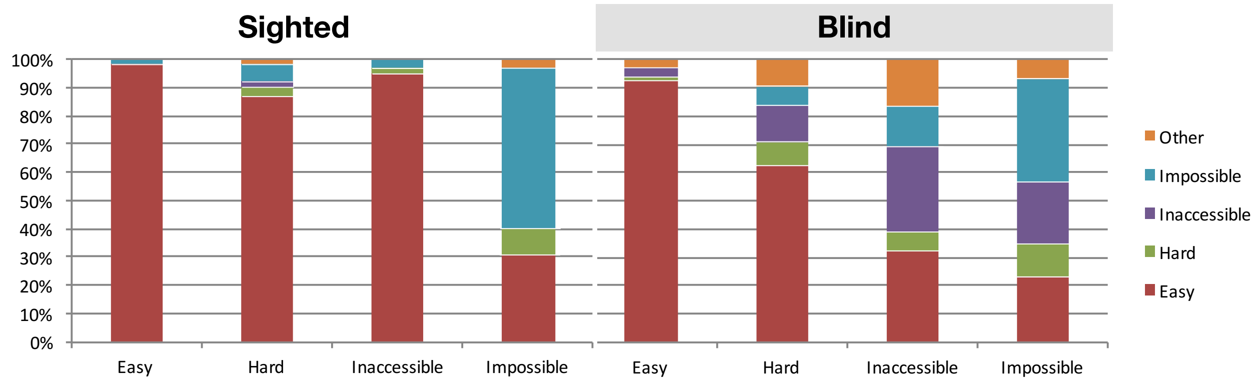

Figure 4: Distribution of blind and sighted participant ratings by difficulty. Sighted participants did not experience a change in task difficulty, and therefore tended to label all of the Easy, Hard, and Inaccessible tasks as being easy. In contrast, as the tasks got more difficult, blind participants indicated they had more difficulty, and frequently confused Hard, Inaccessible, and Impossible tasks.

lems in our tasks most easily, they may provide clues into how to make such tasks easier in the future.

Power users revealed themselves on tasks of each difficulty. On Easy tasks, they described their process of finding the requested information. For example:

> *"If you use the screenreaders ability to bring up links, move to sign up, and below that is the twitter information."*
> *– Task5 (EASY)*

Despite employing various clever strategies, power users sometimes had difficulty with the Hard tasks, such as in this example:

> *"I first looked through the tabbed items (links, buttons, form controls etc) that the site offered and came across mostly buttons. Links seemed to only exist for the contact information. Second, I searched for the word "award" using my screenreader's find facility. It mentioned several times "most awarded" and then a single word on its own line: "Awards". I thought I had hit my jackpot until I saw the list below and saw text quoted from magazines and the like. As a last resort, I even tried looking through the HTML source and could see only a jumble of div, span and path tags with the text seemingly buried way beneath this complete haystack."*
> *– Task2 (Hard)*

The following two example give hints as to how the power users among our participants were able to complete "inaccessible" tasks:

> *"Read the menu. I was interested in all they might have. At the bottom, I found "indicates spicy" I read character at a time to see if there was a symbol my screen reader did not announce. There wasn't, so I checked the font and found spicy was in red. Read the $8 noodles again checking the font of each line. Found the first red one."*
> *– Task4 (Inaccessible)*

> *"I used my favorite method, which was the JAWS find command. I then checked how many times rise was found on the page. Then I checked to see if the placement and time I found for sun rise on the page made sense."*
> *– Task7 (Inaccessible)*

Power users also seemed to choose when to give up based on how much they cared about receiving the information:

> *"Screen reader find function (search for 'clients'). Screen reader mouse cursor function (to try and activate the unlabelled buttons). It's worth noting that, if some of the untagged images on the web page contain client names, I could possibly have obtained them using OCR. I did not do this in this case, because in every day life I wouldn't be aware beforehand that I was looking for client information, so I would simply assume that the information was not there."*
> *– Task2 (Impossible)*

Another power user revealed insight about a realistic mindset when searching for information. They revealed strategies they could have employed, but probably would not use in practice:

> *"I was going to open the page source to find out what <img> or class was used to indicate 'spicy.' I would do this if information was critical to me but on a page like this I would typically live with the mystery."*
> *– Task4 (Impossible)*

6. DISCUSSION

We have explored the effects of NKWYDK on blind web users performing information finding tasks. Overall, our results demonstrated that this does indeed cause confusion, although participants illustrated that the effects of task difficulty are not necessarily straightforward. It was somewhat difficult, for instance, to tease apart the effects of NKWYDK from routine web accessibility and usability problems. Some blind participants also exhibited sophisticated strategies for finding out the information they did not know, and seemed to use heuristics developed through experience to give them insight into what was available on the page (or not). Overall, we believe our study confirmed the impact that subtle usability problems like NKWYDK can have on the experienced accessibility of the web and highlights the richness of this area for future inquiry.

6.1 Sighted vs. Blind Web Browsing

Our study validated differences in sighted and blind web browsing, which were noted at least 10 years ago [6]. For instance, blind people still seem to require about twice as long as sighted people do to accomplish the same task. Interestingly, this same ratio held even when blind people were looking for information that was not contained on the task page, suggesting that the potential added confusion did not slow them down additionally. This inefficiency penalty paid by blind web users certainly has effects on what they are able to accomplish and how effectively they are able to compete. Power users among our blind participants seemed to be especially aware of how the accessibility of web pages could impact their experienced usability of them, reflecting prior results on the relationship between accessibility and usability [22]. These participants occasionally indicated that they did not expect to find the information they were looking for given features of the page and so chose not to try especially hard.

6.2 Difficult, Inaccessible, or Impossible

Participants had nearly as much trouble completing the tasks that we designed to be Hard as they did completing the Inaccessible tasks. This result fits with prior arguments about poor usability being an under-appreciated cause of inaccessibility in some cases [5, 22, 24]. The fact that only about a third of blind participants were able to find information that more than 80% of sighted participants were able to find highlights this continuing problem. Especially novice users may confuse not only inaccessible and impossible tasks, but also difficult, inaccessible, and impossible tasks. This broadened our idea of what kinds of problems fit into the NKWYDK framework to include Hard tasks.

6.3 Confusing Inaccessible with Impossible

Inaccessible information was confused with impossible information in those cases in which the blind participants were unable to get past the accessibility problem. In the individual task feedback, some blind participants reported inaccessible tasks as being impossible, and impossible tasks as being inaccessible. Impossible tasks required nearly an additional minute of time for blind participants to complete. This observed difference suggests that blind web users may be wasting significant time looking for information that is not even there. One of the biggest practical challenges for equalizing access between blind and sighted people is figuring out how to lower the time penalty blind users pay. Future work may consider interfaces that could inform blind web users of the likelihood of finding a particular piece of information given the patterns of what data is usually shared in other websites and also the accessibility that was detected on the page.

6.4 Limitations

Like most studies, ours had limitations. The blind population that we recruited was likely more technically savvy than the average blind computer user, and certainly more technically savvy than the average blind person (using a screen reader is not easy, and so those who master it can generally be considered technically adept). They used a variety of different screen readers, had varying levels of experience with those screen readers, and various levels of experience with the Web. Some were able to solve the tasks that were technically inaccessible, although we were able to separate out these tasks in our analysis. Given that participants likely had no intrinsic motivation to answer the questions, they might have given up more easily than they would have if they truly needed the information they were looking for. Despite these limitations, which mostly leaned toward users we would expect to be best prepared to handle confusing situations, we still observed effects of usability causing substantial inefficiency for our blind participants and NKWYDK appearing to cause substantial confusion.

7. FUTURE WORK

In this paper, we have shown that not knowing what you don't know about information availability and accessibility can affect how successful blind users are in browsing the web. Our findings suggest a number of opportunities for future work in better understanding the effect of NKWYDK on web users and in creating tools to help mitigate its effects.

The first is in expanding our findings in web accessibility. Our study was with blind web users on a specific set of designed web pages; it would be interesting to validate these results in a larger-scale study *in-the-wild*. It is likely that other users may also experience problems from not knowing what they don't know. For instance, deaf web users miss out on auditory information, and web users who do not use the mouse miss out on content that is only accessible that way. In some cases, not being able to access this content may cause similar confusion about what is available and how to best access it. We suspect that the kind of information missing and how it affects different user groups is not the same as how it affects blind web users.

We believe our results may generalize to other challenges faced by people with disabilities off of the web. While we highlighted some examples in our related work section, there are likely others. For instance, blind travelers likely face similar problems when navigating unknown physical spaces (especially indoors where navigation assistance is much more limited), or realizing when their environment changes [2]. There may be an opportunity to develop a uniting theory of NKWYDK that would connect these challenges, and potentially even predict when they are likely to occur.

Finally, we believe there are opportunities to develop tools to help users overcome the problems that we have described here, and that this represents a rich area for future research. We noticed that our more technically savvy blind participants used their knowledge of HTML and CSS in order to access inaccessible information. One idea is to make it easier for less technically savvy users to employ some of the same tactics. Tools might also somehow expose accessibility information to web users so that they can make better decisions about how much time to invest looking for information given global information about a particular site, that either may not contain that information or where the information they are looking for is actually inaccessible. Prior approaches for automatic assessment [30], machine learning prediction of accessibility [3], and end user approaches for fixing inaccessible content [8] may be relevant.

8. CONCLUSION

In this paper, we have explored the concept of not knowing what you don't know, and how it affects blind web users as they complete tasks on the web. A specific and actionable finding from this paper for future research is that confusion between not knowing if information exists on a web page or is simply inaccessible leads to confusion and wasted time, especially for blind web users. Information finding is still less efficient for blind users than it is for sighted users. Power users have developed strategies requiring significant skill to overcome some of these effects, but these strategies are not easily available to less technical users. Our results demonstrate opportunities for future research in further understanding this phenomenon and creating tools to help address it.

9. ACKNOWLEDGEMENTS

This work was developed under grants from the National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR grant numbers 90DP0061 and H133E140039), as well as research grants from the J. Wayne and Kathy Richards Faculty Fellowship in Engineering. We thank Sina Bahram for comments on our study design, and thank the participants in our studies for their time and feedback.

10. REFERENCES

[1] Hervé Abdi. 2007. The Bonferonni and Šidák corrections for multiple comparisons. *Encyclopedia of measurement and statistics* 3 (2007), 103–107.

[2] Nikola Banovic, Rachel L. Franz, Khai N. Truong, Jennifer Mankoff, and Anind K. Dey. 2013. Uncovering Information Needs for Independent Spatial Learning for Users Who Are Visually Impaired. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '13)*. Article 24, 8 pages.

[3] Jeffrey P. Bigham 2007. Increasing web accessibility by automatically judging alternative text quality. In *Proceedings of the 12th international conference on Intelligent user interfaces IUI '07*. 349–352.

[4] Jeffrey P. Bigham, Jeremy T. Brudvik, and Bernie Zhang. 2010. Accessibility by Demonstration: Enabling End Users to Guide Developers to Web Accessibility Solutions. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '10)*. 35–42.

[5] Jeffrey P. Bigham and Anna C. Cavender. 2009. Evaluating Existing Audio CAPTCHAs and an Interface Optimized for Non-visual Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. 1829–1838.

[6] Jeffrey P. Bigham, Anna C. Cavender, Jeremy T. Brudvik, Jacob O. Wobbrock, and Richard E. Ladner. 2007. WebinSitu: A Comparative Analysis of Blind and Sighted Browsing Behavior. In *Proceedings of the 9th International ACM SIGACCESS Conference on Computers and Accessibility (Assets '07)*. 51–58.

[7] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samual White, and Tom Yeh. 2010. VizWiz: Nearly Real-time Answers to Visual Questions. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. 333–342.

[8] Jeffrey P. Bigham and Richard E. Ladner. 2007. Accessmonkey: A Collaborative Scripting Framework for Web Users and Developers. In *Proceedings of the 2007 International Cross-disciplinary Conference on Web Accessibility (W4A) (W4A '07)*. 25–34.

[9] Yevgen Borodin, Jeffrey P. Bigham, Glenn Dausch, and I. V. Ramakrishnan. 2010. More Than Meets the Eye: A Survey of Screen-reader Browsing Strategies. In *Proceedings of the 2010 International Cross Disciplinary Conference on Web Accessibility (W4A) (W4A '10)*. 10 pages.

[10] Erin Brady and Jeffrey P. Bigham. 2014. How Companies Engage Customers Around Accessibility on Social Media. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility (ASSETS '14)*. 51–58.

[11] Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P. Bigham. 2013. Visual Challenges in the Everyday Lives of Blind People. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. 2117–2126.

[12] Giorgio Brajnik, Yeliz Yesilada, and Simon Harper. 2010. Testability and Validity of WCAG 2.0: The Expertise Effect. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '10)*. 43–50.

[13] Ben Caldwell, Michael Cooper, L Guarino Reid, and Gregg Vanderheiden. 2008. Web content accessibility guidelines (WCAG) 2.0. *WWW Consortium (W3C)* (2008).

[14] Ed H Chi, Peter Pirolli, Kim Chen, and James Pitkow. 2001. Using information scent to model user information needs and actions and the Web. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 490–497.

[15] W Vanderheiden Chisholm and G Jacobs. 1999. I.(editors)(1999) Web Content Accessibility Guidelines (WCAG). *W3C* (1999).

[16] Juliet M Corbin and Anselm Strauss. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative sociology* 13, 1 (1990), 3–21.

[17] Vicki L. Hanson and John T. Richards. 2013. Progress on Website Accessibility? *ACM Trans. Web* 7, 1, Article 2 (March 2013), 30 pages.

[18] Jonathan Lazar, Aaron Allen, Jason Kleinman, and Chris Malarkey. 2007. What Frustrates Screen Reader Users on the Web: A Study of 100 Blind Users. *International Journal of Human-Computer Interaction* 22, 3 (2007), 247–269.

[19] Jennifer Mankoff, Holly Fait, and Tu Tran. 2005. Is your web page accessible?: a comparative study of methods for assessing web page accessibility for the blind. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 41–50.

[20] Tara Matthews, Scott Carter, Carol Pai, Janette Fong, and Jennifer Mankoff. 2006. *Scribe4Me: Evaluating a Mobile Sound Transcription Tool for the Deaf*. Springer Berlin Heidelberg, Berlin, Heidelberg, 159–176.

[21] Helen Petrie, Fraser Hamilton, Neil King, and Pete Pavan. 2006. Remote Usability Evaluations With Disabled People. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. 1133–1141.

[22] Helen Petrie and Omar Kheir. 2007. The Relationship Between Accessibility and Usability of Websites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. 397–406.

[23] Peter Pirolli and Stuart Card. 1999. Information foraging. *Psychological review* 106, 4 (1999), 643.

[24] Christopher Power, André Freire, Helen Petrie, and David Swallow. 2012. Guidelines Are Only Half of the Story: Accessibility Problems Encountered by Blind Users on the Web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. 433–442.

[25] Ulrike Schultze and Charles Stabell. 2004. Knowing What You Don‘t Know? Discourses and Contradictions in Knowledge Management Research. *Journal of Management Studies* 41, 4 (2004), 549–573.

[26] Hironobu Takagi, Chieko Asakawa, Kentarou Fukuda, and Junji Maeda. 2004. Accessibility designer: visualizing usability for the blind. In *ACM SIGACCESS Accessibility and Computing*. 177–184.

[27] 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)*.

[28] Hironobu Takagi, Shinya Kawanaka, Masatomo Kobayashi, Daisuke Sato, and Chieko Asakawa. 2009. Collaborative Web Accessibility Improvement: Challenges and Possibilities. In *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility (Assets '09)*.

[29] Jim Thatcher, Cynthia Waddell, and Michael Burks. 2002. *Constructing accessible web sites*. Vol. 34. Glasshaus Birmingham.

[30] Markel Vigo and Giorgio Brajnik. 2011. *Automatic web accessibility metrics: Where we are and where we can go*. Interacting with Computers, Vol. 23:2: 137–155. Oxford University Press Oxford, UK.