

Analyzing the (In)Accessibility of Online Advertisements

Christina Yeung
University of Washington
Seattle, Washington, USA
cyeung3@cs.washington.edu

Tadayoshi Kohno
University of Washington
Seattle, Washington, USA
yoshi@cs.washington.edu

Franziska Roesner
University of Washington
Seattle, Washington, USA
franzi@cs.washington.edu

ABSTRACT

Ads are often designed visually, with images and videos conveying information. In this work, we study the accessibility of ads on the web to users of screen readers. We approach this in two ways: first, we conducted a measurement and analysis of 90 websites over a month, collecting ads and auditing their behavior against a subset of best practices established by the Web Content Accessibility Guidelines (WCAG). Then, to put our measurement findings in context, we interviewed 13 blind participants who navigate the web with a screen reader to understand their experiences with (in)accessible ads. We find that the overall web ad ecosystem is fairly inaccessible in multiple ways: many images are missing alt-text, unlabeled links make it confusing for folks to navigate, and closing ads can be tricky. But, there are straightforward ways to improve: because only a few large companies dominate the ad ecosystem, making small changes to the way they enforce accessibility standards can make a large difference.

CCS CONCEPTS

• Information systems → Online advertising; • General and reference → Measurement; • Human-centered computing → User studies; • Social and professional topics → People with disabilities.

KEYWORDS

Online advertising; measurement; user studies

ACM Reference Format:

Christina Yeung, Tadayoshi Kohno, and Franziska Roesner. 2024. Analyzing the (In)Accessibility of Online Advertisements. In *Proceedings of the 2024 ACM Internet Measurement Conference (IMC '24), November 4–6, 2024, Madrid, Spain*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3646547.3688427>

1 INTRODUCTION

Online advertisements are everywhere: one marketing report from Statista [1] estimates that advertisers will spend almost 300 billion dollars in the US in 2024. But, ads are often designed visually, with images and videos used to convey information. Because most ads are created with a sighted audience in mind, they are not always designed in accessible ways for those who use a screen reader to navigate the web. Prior work has highlighted instances of inaccessible behavior: in 2001, Thompson and Wassmuth [33] found



This work is licensed under a Creative Commons Attribution International 4.0 License.

that half of the ads on news websites were missing alt-text, while Kodandaram et al. [21] surface dark patterns in online ads that make it hard for people who use screen readers to know when they are interacting with ads. Additionally, anecdotal reports from the news have also shown that the designs on large social media platforms can lead to inaccessible behaviors: Facebook developers used HTML such that ads were not labeled as third-party content for blind users for two years [26].

Motivated by these older and more recent anecdotal concerns, our overarching questions are: (1) What are the accessibility practices of ads on the web today? (2) What are the experiences of screen reader users with these ads?

We break accessibility for ads down into three components, using a subset of best practices from the Web Content Accessibility Guidelines (WCAG). First, we examine whether ads are *perceivable*: How do ads present information to users, with particular regards to assistive attributes such as ARIA-labels and alt-text? Second, are they *understandable*: Does a screen reader user know that they are interacting with an ad, and are they able to understand what the ad is trying to promote? Finally, are they *navigable*: can someone using a screen reader easily make decisions about whether or not they want to interact with advertisements when they encounter them? Or, do online advertisements disrupt the way they would otherwise be able to browse the internet?

We use a mixed-method approach to answer these questions. First, we run a month-long web measurement study of ads on 90 popular websites from 6 different categories. We use the HTML of these ads to audit their accessibility characteristics, using a subset of the WCAG 2.2 guidelines, including checking for missing alt text, missing text associated with links, and missing text associated with buttons. We additionally create a new category of behavior we believe is inaccessible: ads that only contain “non-descriptive” text, such as “Advertisement,” or “Ad” to describe their content.

Second, to provide context and a deeper understanding of our web measurement results, we conducted semi-structured interviews with 13 blind participants. In the interview, we discussed people’s experiences with and opinions of online ads to understand the designs they found the most accessible, as well as the designs that posed challenges in practice. We also asked participants to navigate through a website that we designed, which hosted several ads we observed during our measurement study, to understand real users’ reactions to the inaccessible characteristics we quantified in our measurement results.

Overall, we find that not only do web ads today have significant accessibility limitations (via our web measurements), but also that those limitations have direct negative impacts on users (via our semi-structured interviews). Fortunately, many of the fixes are technically straightforward, and small changes can make a large impact. We are in the process of reaching out to relevant ad platforms to

share our findings and urge them to improve the accessibility of their ads. As part of our contribution, we have made our data available at <https://ads.cs.washington.edu/projects/adaccessibility.html>.

2 BACKGROUND

2.1 Terminology

We use the term “advertisement” to describe the online content an advertiser has made to promote their product — for example, an ad for Coca-Cola. Advertisers are responsible for the way the advertisement is constructed, including the underlying HTML and JavaScript used to display the ad. Advertisers can include text, or other tags, to easily disclose an ad’s status as third-party content.

We also discuss “advertising platforms”, which are companies that facilitate the delivery of advertisements on the web. These entities work with websites to help place ads when needed. They also have the ability to add infrastructure through HTML and JavaScript around the advertisement, such as information that tells users why they were served ads.

Finally, “websites” are the pages which offer space for advertisements to be placed. Website owners, too, can use HTML to denote the spaces before and after ads are placed.

2.2 WCAG Standards

The Web Content Accessibility Guidelines (WCAG) are a set of internationally recognized standards developed by the World Wide Web Consortium (W3C) [32] to guide web developers, designers, and content creators in making their digital content more accessible to people with disabilities. The latest guidelines, the WCAG 2.2 standards, build on four major principles. In our work, we focus on three of the four principles: perceivability, understandability, and navigability. In this section, we briefly review what each of these principles mean, and in Section 3, we elaborate on how we use them to interpret the accessibility of ads.

Regarding perceivability, WCAG states that accessible web elements must provide alternatives for non-text content. Examples of this include providing alternative text (alt-text) for images, or subtitles for video content. Second, understandability: web elements must make their content clear and readable to users. The third principle broadly is called operability: we focus on a specific component of it called navigability, where components and navigation elements must be functional through keyboard navigation.

The final principle, robustness, states that web content should be interpreted accurately by a variety of user agents and assistive technologies. Though important, we do not analyze the robustness of ads. Future work can expand on our research by examining other assistive technologies, including screen magnifiers, or text-to-speech interfaces.

We emphasize that though these principles are presented individually, they can affect each other. For instance, a web element that contains no perceivable information — that is, it does not expose anything to someone who browses the internet with a keyboard — then by definition, it will also not be navigable. Another example may be content can be perceivable, where screen readers are able to convey information — but there is no guarantee that the information will be understandable. Finally, a web element that is not

understandable (for instance, if an individual cannot understand the content of an ad) is also much harder to navigate.

2.3 Accessibility Tree

In order to study the accessibility of online advertisements, we leveraged the accessibility tree (computed by the browser) to retrieve the information they expose to screen readers. This tree provides a structured representation of the content of a web page that assistive technologies, such as screen readers, can interpret and interact with. It is derived from the Document Object Model (DOM), but specifically focuses on elements that are perceivable and operable by people who use keyboards to navigate websites. Note that, as the tree is derived from the DOM, each browser has its own implementation of the accessibility tree. Thus, an accessibility tree of a page in Chrome, compared to Firefox may differ: nevertheless, the underlying goal of the tree is to provide an accessible experience for people who use screen readers (as well as other assistive devices that use the accessibility tree to convey information, such as braille readers).

The accessibility tree includes information that helps people understand what the HTML element is, such as alt-text for images, or labels for form fields. It contains five pieces of information for each element: the *name* of the element, also known as its “accessible name.” This is the text that a screen reader will announce, when the element is focused on. Depending on the element, the accessible name can be derived in multiple ways, including ARIA-labels (which are labels explicitly to convey descriptions for accessibility purposes), titles, alt-text, and the actual text in the body of the element. The accessibility tree also contains a *description* of the element, which provides more context than just the accessible name. However, depending on the screen reader people use, the descriptive text is not always read out by default — sometimes, screen readers alert users that there is additional descriptive text at the end so that they can choose to interact with it if they desire. The third element of an accessibility tree is the *role* of an element, like whether or not it is a button or a link. The fourth element is the *state* of an element, such as whether or not a checkbox is checked or unchecked. Finally, there is the *focusability* of an element — that is, whether or not it can be accessed via keyboard.

3 METHODOLOGY: MEASUREMENT

3.1 Collecting Ad Dataset

3.1.1 Selecting websites. We collected ads on popular news, health, weather, travel, shopping, and lottery websites in the US, as they commonly embed ads to generate revenue. We selected these by using SimilarWeb [36], a service that aggregates and ranks websites by popularity both by broad category, and by country. We started by selecting the most popular websites in each category, and manually visited them to ensure they served ads. If the site did not appear to deliver ads, we removed it from our list of sites, and moved on to the next-most-popular site in ranking. In total, we selected 90 websites to crawl: the top 15 sites from each category that embedded ads. Note that in the specific case of travel sites, the landing pages didn’t display ads directly: instead, the search results subpages do. As such, for each of the 15 travel sites, we searched for travel between the same two cities, using the same date ranges.

3.1.2 Crawling and collecting ads. To visit websites and collect ads, we use AdScraper [40], a tool that uses Puppeteer to navigate to URLs and save the ads observed on each site. When visiting a page, it closes out of any pop-ups, scrolls up and down, and identifies ad elements using EasyList CSS rules. AdScraper saves a screenshot of each ad element, as well as its HTML content. When ads are served in nested iframes, AdScraper iterates through each level to get to the innermost available HTML. We additionally modify AdScraper to capture each ad's accessibility tree by using the Chrome DevTools Protocol API: this allows us to parse the information it exposes to screen readers in an automated way. We visited each URL with a clean profile and cleared cookies between each page visit.

3.1.3 Post-processing collected ads. When manually examining our collected ads, we observed instances where we were unable to fully capture the ad images for two reasons: first, where the ad was not able to fully load before we captured it, and second, where we were unable to capture the full HTML of ads. Both of these instances resulted in ad screenshots that were composed of only whitespace, and its saved HTML was incomplete. As such, we processed ads that we collected by examining the pixels in each ad's screenshot. If all the pixels in a screenshot had the same value, we classified it as one that was a blank screenshot. We also checked each ad's saved HTML, using a parser to determine if the content began and ended with the same tag: if it did not, we categorized it as incomplete. If an ad did not pass either check, we removed the ad from our final data set. This may have occurred due to the dynamic nature of ad delivery: in all of these instances, the scraper identified a region as an ad, but prior to scraping its content, a different ad was delivered in its place.

We also deduplicated the ads, using an average hashing function, as well as the contents of their accessibility tree. Note that we used both an ad's image, as well as the content it exposed to screen readers when deduplicating, particularly because ads that visually look the same might not share the same information to assistive devices.

3.1.4 Final data set. For our final data, we analyzed ads that were collected over the course of a month, from January 20, 2024 to February 21 2024. In total, we collected 17,221 total ads over the course of the 31 days. After deduplication, these 17,221 total ad impressions correspond to 8,338 unique ads. After post-processing to ensure that we were able to successfully capture the HTML for each ad, our final data set included 8,097 unique ads. We will make our dataset of ads, accessibility tree data, and analysis code publicly available.

3.1.5 Identifying ad platforms. In order to identify ad platforms that deliver ads, we manually looked for visual heuristics that characterize ad platforms, and identify the corresponding HTML elements. We used two main heuristics: first, the “AdChoices” button present on some ads to explain why the ad was delivered to the individual. After identifying ads with this button, we examined where clicking on the button would lead, by inspecting the HTML to extract the relevant URL.

The second visual heuristic we used was when platforms displayed their name alongside the ad. For instance, many native ads are presented in grids of thumbnails, with a visual indicator at the

top saying “Ads by [COMPANY NAME].” We also extracted the URLs associated with these platforms. This is an iterative process: once we identified one characteristic, we applied the label to our data set, and analyzed the HTML of ads that had not yet been identified as delivered by a specific ad platform. If the ad's HTML contains a platform's URL, we identified it as an ad that was delivered by that platform, and performed an additional manual check by looking through the ad's HTML to make sure it truly included the ad platform's visual heuristics. In total, we manually analyzed 2,000 images, and identified URLs associated with 16 ad platforms.

From there, we applied these heuristics to the 8,097 unique ads in our data set. We found that we were able to identify the ad platforms that delivered 5,817 (71.9%) of the ads. We restrict our current analysis to platforms who delivered at least 100 unique ads in our dataset: this includes 8 advertising platforms, which collectively delivered 71% of the unique ads in our data set. These 8 platforms are: Google, Taboola, Yahoo, Criteo, The Trade Desk, Amazon, Media.net, and OutBrain.

3.2 Analyzing the Inaccessibility of Ads

In this section, we explain the methods we used to determine whether or not an ad contained inaccessible content. We use three principles from WCAG — perceivability, understandability, and navigability — as a way of assessing the behaviors we observed in our collected data. If an ad was not accessible in any of the three dimensions, we classified it as an ad that exhibited at least one inaccessible characteristic.

3.2.1 Perceivability. When considering the principle of perceivability — the WCAG concept that web content should have textual representation of visual cues — we focus on analyzing two components: first, we look at the types of HTML assistive attributes ad developers use to expose information (such as alt-text, ARIA-labels, and so on). Second, we examined the alt-text property in more detail.

HTML assistive attributes. The first aspect of perceivability that we examine is to measure the prevalence of various HTML assistive attributes that developers use to expose information, such as ARIA-labels, alt-text, titles, and the actual contents of the HTML tag (if it is exposed to screen readers).

Note that two visually identical ads may differ in their accessibility attributes. For instance, consider the example in Figure 1, which shows two different ways the developer of an ad can choose to display a clickable image. In both implementations, a user would be able to click on the image of the flower, and arrive at example.com. However, the two methods diverge when it comes to what information is presented to screen readers. Here, we would consider the bottom (HTML+CSS) implementation to exhibit more inaccessible characteristics — specifically, we would find that it contains more content that is not perceivable to a screen reader — compared to the HTML-only implementation. The reason for this is that the HTML-only version contains alt-text for the image (“White flower”) that is exposed to screen readers. By contrast, the HTML+CSS version (though perhaps better from a web development perspective for variable layouts on different devices) lacks an alt-text property. Depending on the screen reader a person is using, the assistive



```
HTML-only implementation:
<a href="https://example.com">
  
</a>
```

```
HTML+CSS implementation:
<div class="image-container">
  <a href="https://example.com">
    <div class="image"></div>
  </a>
</div>

.image-container { display: inline-block; }
.image {
  width: 300px;
  height: 200px;
  background-image: url('flower.jpg');
  background-size: cover; }
a { text-decoration: none; }
```

Figure 1: This flower represents a clickable image on a webpage, with two possible implementations shown.

technology may try to present more information, such as reading out the contents of the anchor tag (the URL), but this may or may not be text that is understandable.

Alt-text. In addition to counting the presence of different assistive attributes, we also take a deep-dive into alt-text. Images without alt-text are not perceivable by people who use screen readers: though their assistive technology may inform them that an unlabeled image is present, they do not have a straightforward way of explaining what is contained in the image without using additional third party software.

In order to examine if visible images in each ad were missing alt-text, we analyzed each ad for the HTML image tag, and ignored any images that are smaller than 2x2. We also did not include images whose CSS display or visible attributes are set to ‘none,’ or ‘hidden.’ For all other images, we check to see if each image in an ad has a corresponding alt-text. If the alt-text property is not present (e.g., ``) or if it is present, but only contains an empty string (e.g., ``), we consider the ad as containing an image with missing alt-text. Note that though in the latter case, some developers use this empty string as a way to denote “decorative” images that should be skipped, some screen readers will still perceive the content as an image that is missing alt-text. As such, in this work, we still consider it behavior that is not perceivable by a screen reader.

Table 1: Strings denoting ad disclosure

Word	Suffixes
ad	-s, -vertiser, -vertising, -vertisement, -vertisements
sponsor	-s, -ed, -ing
promot	-e, -ed, -ion, ions
recommend	-s, -ed
paid	N/A

3.2.2 Understandability. In our work, we consider three elements when analyzing the understandability of content: first, we look at whether the ad discloses its status as third-party content. Second, we analyze whether the content only contains what we term “non-descriptive” text. Finally, we consider ads that have no associated text with their link.

Ad disclosure. When analyzing the understandability of ads, we first examined whether or not ads disclose their status as third-party content to screen readers. This is, in part, motivated by the prior work that shows that platforms used designs that made it hard for blind folks to distinguish between the organic content on a page, and the content of an ad [26]. Further, the Federal Trade Commission’s guidance on .com Disclosures [14] states that clear and conspicuous disclosure of online advertising helps consumers make informed decisions. While the language in the disclosure alone might not prevent misunderstanding by itself, it is a key component in the way people can accurately interpret the claims made in ads. And, we believe that appropriate disclosure is one way ads can inform people who use screen readers that they are, indeed, interacting with an ad.

To determine whether ads in our data set contain strings that disclose their status as ads, we first split our data set in half. For half of the unique ads in our data set, we manually examined the content exposed in each ad’s accessibility. If an ad contains language that disclosed its status as third party content, we extracted the term that was responsible for disclosing its status. Examples of strings include the word “Advertisement.” We then deduplicated the list of words that ads use to disclose their status, present the set of words and their corresponding suffixes in Table 1.

After manually reviewing half of our total data set of unique ads, and identifying the terms used to disclose third-party status, we searched through the remaining, unlabeled half of ads’ accessibility tree. If any element in the ad, including links, images, buttons, or text, contained any of the keywords denoting third party status, we considered the ad as having disclosed its status.

Finally, after checking all the ads in our data set for whether any piece of information contained a disclosure, we manually reviewed each of the ads that were not labeled as such. That is, if an ad was categorized as lacking a disclosure, we manually inspected the ad to check for the presence of previously undetected language that might denote its third-party status – we did not discover any additional language, and did not change the labels on any ads. As such, the ads labeled as lacking disclosure in our data set truly represent ads that do not indicate, to a screen reader, that they are ads.

“Non-descriptive” text. We next turned our attention to analyzing all of the information an ad exposes to screen readers. As mentioned in a prior section, information can be perceivable by a screen reader – it might detect that there are links in an ad – but this information may not be understandable to the user. To examine the information present in ads, we examined all of the unique ads in our data set. We separated the strings of each ad by the type of attribute (e.g., alt-text, titles) individually.

For each type of attribute, we deduplicated the strings observed from across all ads, and manually reviewed the unique strings, sorted from most frequently observed to least. As we did this, we labeled each string as either “non-descriptive,” or “contained text specific to an ad.” Examples of “non-descriptive” text might be an ARIA-label that says “Advertisement,” a title that says “3rd party ad content,” or an alt-text for an image that says “Image.”

Table 2 shows the three most common strings for each assistive attribute across our dataset, as well as the count of unique ads that used that particular language in their disclosure. After deduplicating the strings from each attribute, we tokenized the “non-descriptive” strings, and counted the number of ads whose attribute only contained those words. If the content of ads *only* contained generic words, we labeled the whole ad as non-descriptive. For example, a non-descriptive ad in our data set might expose information that says “Advertisement” from the ARIA-label of an iframe, followed by “Learn More” from the contents of a link.

Missing (or generic) text associated with links. The last element we assess the understandability of advertisements by looking at the number of ads that are missing text associated with the links in their content. The text associated with links inform users what will happen if they click on it. For instance, an HTML for a link with associated text might look like the following:

```
<a href="http://example.com/">Example text that gets
conveyed to users </a>.
```

In this example, a person who uses a screen reader would see this link, and receive the text contained within the tag (that starts with “Example” and ends with “users”).

We consider two types of text associated with links to be not understandable: first, instances where the text is non-descriptive. Text that includes content such as “learn more” does not help orient users, or provide information as to what might happen when they click on the link.

Second, we consider instances where there is no text associated with a link. In contrast to a link with associated text, an empty hyperlink does not have the text within the `<a>` tag, and may follow a similar format as the following:

```
<a href="http://example.com/"></a>.
```

People who use screen readers will be able to navigate to the link, because link elements get keyboard focus by default, but may not be able to understand what the content actually means. Some screen readers that encounter a hyperlink with no associated text say “link,” which we consider non-understandable behavior, as it is non-descriptive.

Other screen readers that encounter a link with no associated text may start reading the contents of the href out letter by letter. The nature of advertisements can make this particular behavior especially difficult to understand: oftentimes, the URL presented

to the user is not the domain of the final landing page, but rather, an intermediary used for click attribution purposes. For example, some ads delivered by Google use its advertising company’s domain: `doubleclick.com`, followed by a series of numbers and strings for attribution purposes. Doubleclick may not be a familiar domain to all: not everyone would understand what that domain means, or what the landing page ultimately will resolve to. Compounding challenges to understanding, the series of numbers and strings does not often hold significant meaning for people – it is a way of attribution so that the advertiser can gain insights into how people arrived at their landing page.

3.2.3 Navigability. Navigability represents the idea that web users must be able to interact with, and browse through content effectively. When assessing the navigability of ads, we focus on two aspects: first, the number of interactive elements an ad contains. And, second, whether the ad contains buttons that have missing or non-descriptive text.

Number of interactive elements. When considering the number of interactive elements, we use the accessibility tree of each ad, and examine how many elements can be discovered as someone presses the tab key to traverse through the content. Note that this number is likely a lower bound, as ads might contain more content that is not necessarily keyboard navigable, such as text in divs and spans that don’t inherently have tab focus: instead, users would need to use arrow keys or other shortcuts to access this information.

We consider ads that contain 15 or more interactive elements as content that is not navigable. This means that a user who traverses content linearly (i.e., uses the tab key to navigate through an ad) needs to press the tab key 15 times to reach other content on the website.

Missing text associated with buttons. Our last metric for considering whether ads contain inaccessible content is through checking the text associated with the buttons that they may contain. Buttons, similar to links, receive keyboard focus by default, and expose the text contained within the `<button>` tags to screen readers. They often affect a user’s ability to navigate, as buttons are commonly used to close out of ads.

We consider an ad non-navigable if it contains buttons that do not have any associated text. In this case, screen readers will still navigate to the button, but instead of informing users what they might be able to do, like exit out of the ad, the screen reader will announce the word “button.” Without text describing the functionality of the button, screen reader users cannot differentiate buttons that will click the ad, close out of the ad, or provide more information about the ad.

4 RESULTS: MEASUREMENT

We present the results of our measurement of the ads we observed over 31 days of scraping 90 websites. We summarize high-level findings in Table 3 and break down what each row means in the following sections in more detail. Overall, we find that only 13.2% of the ads in our data set (1,069 unique ads) do not exhibit any inaccessible characteristics.

Note that each ad can exhibit more than one type of inaccessible behavior – that is, it can be missing alt-text (not perceivable),

Table 2: Most commonly observed strings for each assistive attribute

ARIA-label	Title	Alt-text	Contents
Advertisement (3,640)	3rd party ad content (3,640)	Advertisement (697)	Learn more (1,603)
Sponsored ad (345)	Advertisement (914)	Ad image (20)	Advertisement (837)
Advertising unit (42)	Blank (90)	Placeholder (20)	Ad (411)

Table 3: Inaccessible Characteristics of Ads

	Count	Percentage of all ads	Type of inaccessible behavior
Has no alt, empty alt string, or non-descriptive alt	4,600	56.80%	Perceivability
Ad does not contain disclosure	511	6.30%	Understandability
Information is all non-descriptive	2,838	35.10%	Understandability
Missing, or non-descriptive link	5,057	62.50%	Understandability
Ads with ≥ 15 interactive elements	202	2.50%	Navigability
Missing text for button	2,476	30.6%	Navigability
<i>Ads without any inaccessible behavior</i>	<i>1,069</i>	<i>13.2%</i>	<i>None</i>

Table 4: Accessibility of Ad Attributes

	Total	Non-descriptive or empty strings	Contained text specific to ad
ARIA-label	5,725	5,026 (87.8%)	699 (12.2%)
Title	8,010	6,805 (85%)	1,205 (15%)
Alt-text	5,251	3,267 (62.2%)	1,984 (37.8%)
Tag contents	45,436	15,037 (33%)	30,399 (67%)

have missing text associated with its link (not understandable), and be composed of more than 15 interactive elements (not navigable). Note also that we count the number of ads that exhibit each inaccessible characteristic, not the total number of times each characteristic occurs across our data set (i.e., an ad with two images lacking alt-text is only counted once in the first row of Table 3).

4.1 Perceivability

4.1.1 Most assistive attributes contain non-descriptive language. As outlined in Section 3.2.1, ad developers have multiple ways of exposing ad content to screen readers. Table 4 shows the different types of assistive attributes ad developers use to expose the information, separated by whether or not the information was non-descriptive. We find that all of the ads in our data set expose at least one piece of information to screen readers, in one of four ways: ARIA-labels, titles, alt-text, or directly exposing the information contained within HTML tags. We found that ads commonly separate the information they present to screen readers into multiple parts, so an ad that visually looked as though it may be one unit could have up to 40 components in different HTML tags.

Unfortunately, we found that the information in assistive attributes contains non-descriptive language, such as “ad” or “image”, more than half the time. More specifically, ARIA labels contain non-descriptive language 87.8% of the time they are used, titles 85% of the time, and alt-text 62.2%.

4.1.2 Over half of ads are missing alt-text entirely, or contain empty or non-descriptive strings. We emphasize our finding about alt-text, which is intended to convey to screen reader users the contents of an image: over half (56.8%) of the ads in our data set contained

text that was either empty, or had non-descriptive text. This breaks down to 26% of ads with no alt-text and 30.8% of ads with non-descriptive alt-text. Though alt-text was first proposed around 1993, its adoption has been slow. Our result is slightly better than a 2001 study [33], which found among banner ads on news websites, 74.73% contained either empty or unhelpful alt-text; this may suggest that alt-text adoption has improved over time.

We note that in ads, alt-text may sometimes be redundant, such as when there is both an image for a logo, as well as the company’s name in text elsewhere in the ad. As such, users’ preferences may vary: under the principle that all visual information should have textual counterparts, some may find it helpful to have alt-text available all the time, even if it provides redundant information. It may also describe additional details about the logo that are not available otherwise, such as the color or shape of the logo itself. On the other hand, some users may find the redundant information annoying.

4.1.3 Ad developers still use titles to convey information, contrary to guidelines. Finally, we consider ad developers’ use of the title attribute. The title attribute allows web content developers (not just ad developers) to provide more context for specific HTML elements. When applied, the title tag is primarily shown as a tooltip, displayed when a user’s mouse hovers over markup elements.

However, relying on the title attribute for accessibility can be problematic, as not all users are able to consistently interact with it. Depending on the screen reader, different assistive technologies will either skip the content in titles entirely, or only in very specific cases convey the information to users. Making matters worse, outdated Search Engine Optimization (SEO) advice led to poor use-cases of title tags [2]. Thus, web accessibility experts advise not to use the title tag at all when trying to expose information to screen readers and other assistive technology [6, 15, 28]; the WC3 discourages relying on it as an accessibility attribute as well [37].

Nevertheless, in our data, we observe several instances where ad developers are still using only the title tag to convey information specific to advertisements. This represents information that will not be perceivable by all users of screen readers. Of the 1,205 unique ads that use the title attribute to convey information specific to an ad, shown in Table 4, many of them do not repeat this information

Table 5: Ad Disclosure Types and Counts

Ad Disclosure Type	Count
Disclosed through keyboard focusable elements	6,063
Disclosed through static text (not keyboard focusable)	1,523
Not disclosed	511

in other ways that expose the information more directly to screen readers.

4.2 Understandability

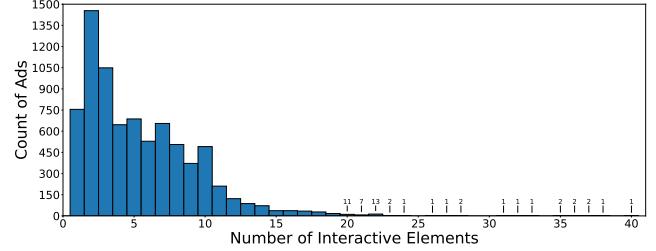
4.2.1 The vast majority ads inform users that they are ads through language that screen readers can detect. While ad disclosures alone do not validate the claims made in an ad, its clear and conspicuous presence may help consumers make informed decisions. We focus on three different ways ads could disclose their status through text: (1) through an element with tab focusability, such as a link; (2) through an element of the ad that discloses, but does not receive keyboard focus by default, such as the text in a div or span tag; or (3) not at all. We separate the first two conditions because disclosures via non-keyboard focusable elements may be missed by people who traverse content quickly.

Table 5 shows the breakdown of disclosure by ads. Each unique ad appears once in this table: we count the first time we observe a disclosure, though an ad could contain multiple. Overall, we find that the vast majority of ads in our dataset disclose their status as ads through text that screen readers can detect. Across disclosures between elements that receive keyboard focus, as well as elements that do not, 93.7% of ads in our data set (7,586 ads in total) identify themselves as ads through text. Thus, our findings suggest that at the very least, ads on the pages we crawled informed users that they were interacting with ads (and not the actual content of the site) through the language exposed to screen readers.

4.2.2 Over a third of ads contain non-descriptive language. While we already reported results on non-description strings in Section 4.1.1, we reinterpret them here in the context of understandability. Though ads with non-descriptive strings may technically satisfy accessibility requirements, users are ultimately not able to understand what the ad is trying to promote. For example, someone who uses a screen reader might not be able to tell the difference between a job advertisement whose alt-text says merely “advertisement” and a malicious ad that also uses the same text. Not only do they miss the opportunity to engage with the job advertisement, they might also be more vulnerable to harmful ads that sound similar to benign content.

As shown in Table 3, we find that over a third (35.1%) of the ads in our data set *only* contain non-descriptive information. This means that screen reader users may not be able to distinguish the differences between 1 out of every 3 ads.

4.2.3 Links are either missing text, or only contain non-descriptive text, in over half of ads. Table 3 also shows that over half the ads in our data set are either missing text, or that the text shared with screen readers is entirely non-descriptive (e.g., screen readers might only say “link” or try to start reading the URL specified in the anchor tag). Indeed, links with missing or non-descriptive text represents

**Figure 2: Distribution of number of elements across unique ads****Figure 3: Ad with 27 interactive elements**

the *most common* reason ads fail to be accessible in our data set. Notably, by WCAG guidelines, ads that contain at least one missing link will not meet the minimum standards required to be considered legally accessible. This could mean that these ads, on websites that otherwise comply with accessibility guidelines, might erode the accessibility of the overall content.

4.3 Navigability

4.3.1 Ads with 15 or more interactive elements occur infrequently, but can be challenging to navigate. Figure 2 shows the distribution of the number of interactive elements in each of the ads we observed. Across our full data set, we see a fairly long tail: the fewest number of interactive elements an ad had was 1, while the largest number of interactive elements in an ad was 40. Overall, we found that most ads contained between 2 and 7 interactive elements: the average ad in our data set contained 5.4 keyboard focusable elements.

In our work, we classify ads with 15 or more interactive elements (that is, elements that are able to receive keyboard focus) as ads that are not navigable. As shown in Table 3, this represents a very small number of unique ads we observed: in fact only applies to 2.5% of the unique ads in our data set. Nevertheless, ads with multiple elements potentially represent content that is harder to navigate through with a keyboard than it may be to skip over visually.

For instance, Figure 3 demonstrates an example of an ad that contains 27 interactive elements. In this ad, each of the shoes are contained in its own anchor tag. Because the many links are also unlabeled, it may be especially difficult to navigate this ad. Depending on the screen reader, users may either hear “link” repeated 27 times as they tab through each shoe, or their screen reader may read the non-human-readable URL associated with each element.

4.3.2 Just under a third of ads contain buttons with missing text. As shown in Table 3, roughly 30% of the ads in our data set contain buttons that do not have associated text. Without text describing the functionality of the button, navigation is hard: screen reader

users have to guess what it does based on the surrounding context. Particularly because buttons are often used to close out of ads, people may accidentally click on unlabeled buttons, mistakenly believing that they are exiting out of the ad.

4.4 Findings by Ad Platform

In previous sections, we examined our overall set of ads, and quantified the different types of inaccessible behaviors we observed. Now, we ask: Are certain ad platforms delivering content that is more (or less) accessible?

4.4.1 Inaccessible ads are unequally distributed among ad platforms. Table 6 shows a subset of the accessible behaviors of ads in our data set, separated by the ad platform that delivered it. Overall, we can see that the inaccessibility of ads is not randomly distributed across ad platforms: instead, there are platforms that appear to serve ads that are accessible more frequently than others. Specifically, ads delivered by the two clickbait companies (Taboola and OutBrain) exhibit relatively more accessible behaviors when compared to other advertising platforms. For example, 42.7% of the ads delivered by Taboola, and 81.5% of the ads delivered by OutBrain did not exhibit any of the characteristics we consider inaccessible. We find that the only other platform that delivered ads with relatively accessible ads appeared to be Amazon (23.7%): most other platforms delivered ads that had at least one inaccessible element over 99% of the time.

4.4.2 Clickbait ads are disproportionately more accessible. We hypothesize that Taboola and Outbrain do better in accessibility — they deliver ads that only contain non-descriptive text less than 1% of the time — because they use relatively standard, HTML-based templates for their clickbait style ads (sometimes called “chumboxes” [24]). Indeed, Zeng et al. [41] found that Taboola and OutBrain deliver essentially *only* low-quality clickbait ads designed to draw clicks by using hyperbolic language, alarming images, or exaggerated claims. (This is not to say that the other platforms never deliver clickbait, but they also deliver substantial numbers of higher-quality ads.) Our finding, combined with that from prior work [41], means that there are implications for screen reader users in terms of which types of ad content they are disproportionately exposed to. That is, the content of ads that a screen reader user hears likely skews more towards clickbait than high-quality ads for products, services, public service announcement, or opportunities (e.g., housing, jobs).

4.4.3 Case Studies. Now, we investigate specifically what makes ads from Google, Yahoo, and Criteo less accessible.

Case study: Google’s unlabeled “Why This Ad?” buttons. As seen in Table 6, Google delivers ads that have missing text with their buttons far more often than any other platform: 73.8% of the ads delivered through Google.

As we inspected the Google ads that contained an unlabeled button, we found that all of them had one unifying characteristic: they were all a result of the “Why this ad” / “Ads by Google” button presented in the ad. Figure 4 shows an example of an ad with this button circled in red, and the interface that a user sees when they click on the button.

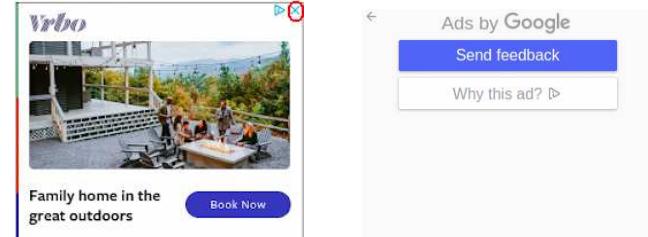
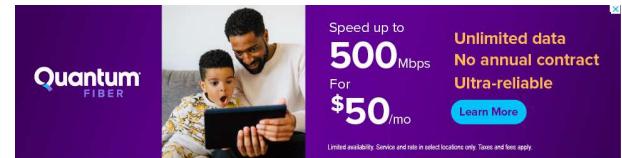


Figure 4: User flow of Google’s “Why this Ad?” process



```
<div style="position: absolute; width: 0px; height: 0px; border: 0px; padding: 0px; margin: 0px; overflow: hidden;">
  <a href="https://yahoo.com"></a>
</div>
```

Figure 5: Example Yahoo ad with an “invisible” div

Ironically, this button is intended to provide more context as to why people see the specific ad that is delivered: however, in the context of content heard by screen reader users, this can actually make an ad less accessible than before. However, the fix is simple: Google needs to update its template such that this label has appropriate language explaining what happens when people interact with the button.

Case study: Yahoo’s visually hidden links. We found that all of the ads Yahoo delivered contained links that either had empty, or non-descriptive text. When we inspected these ads, we found that all of the ads contain a link that is not labeled, is visually hidden, but is still exposed to screen readers. Figure 5 shows an example. The ad contains an unlabeled link leading to yahoo.com nested in a div tag set to 0px. This link is visually hidden but will still be announced by screen readers.

A simple solution would hide this element from screen readers by using additional assistive attributes, such as the ARIA-hidden flag.

Case study: Criteo’s div tags masquerading as (inaccessible) buttons. The main reasons we found for Criteo ads being were empty alt text and no text associated with links. Both issues stem from Criteo’s privacy disclosure and close ad buttons, which are implemented as div tags with CSS to appear as buttons. Figure 6 shows an example of an ad delivered by Criteo, with the privacy and close buttons circled in red, as well as the associated HTML.

Implementing buttons via div tags is discouraged from an accessibility perspective [19], e.g., because div tags do not receive keyboard focus by default, and because, unlike button elements, they lack inherent semantic meaning. The fix is thus simple: Criteo

Table 6: Inaccessible behavior across different platforms

Inaccessible Behavior	Google	Taboola	OutBrain	Yahoo	Criteo	The Trade Desk	Amazon	Media.net
Alt accessibility problems	1,813 (66.5%)	53 (3.2%)	100 (18.5%)	251 (94.4%)	216 (99.5%)	196 (92.9%)	127 (61.4%)	105 (66.5%)
Non-descriptive content	1,344 (49.3%)	4 (0.2%)	0 (0%)	44 (16.5%)	33 (15.2%)	152 (72%)	63 (30.4%)	50 (31.6%)
Missing, or non-descriptive link	1,865 (68.4%)	903 (54.5%)	0 (0%)	266 (100%)	216 (99.5%)	124 (58.8%)	100 (48.3%)	116 (73.4%)
Missing text for button	2,012 (73.8%)	5 (0.3%)	0 (0%)	61 (22.9%)	5 (2.3%)	46 (21.8%)	31 (15%)	47 (29.7%)
Ads without any inaccessible	12 (0.4%)	707 (42.7%)	440 (81.5%)	0 (0%)	0 (0%)	0 (0%)	49 (23.7%)	0 (0%)
Platform total	2,726	1,657	540	266	217	211	207	158

**Figure 6: Example Criteo ad**

could use an ad template in which the button is implemented via the button HTML tag.

5 METHODOLOGY: USER STUDY

Our measurement study surfaced many accessibility issues with ads, based on our interpretation of the WCAG guidelines. However, none of the authors uses a screen reader themselves, and we sought to understand, directly from screen reader users, the actual impact of our findings on their interaction with a perception of ads on the web. Thus, in the second phase of our research, we conducted semi-structured interviews with 13 blind participants to understand how ads affect the way screen reader users browse the web.

Recruitment. We recruited participants in two ways: first, we reached out to organizations that provide services to blind or low vision people in each state. We also posted invitations on social media community groups to invite people to participate in our study. In total, we recruited 1 participant through state contacts, and 12 through social media responses. We paid participants \$30 for their time, and interviews lasted about 45 minutes on average.

Study protocol. Our semi-structured interview protocol consisted of four phases that were conducted in a single, remote session with each participant. The full interview protocol can be found in Appendix A

First, we asked participants about their background in using screen readers and their web browsing habits. Second, we focused on participants' experience with online ads. We asked open-ended questions, focused on two main themes: first, the types of ads people observed, and what they did and didn't like about them, as they browsed the web. Second, we asked people about the way they navigated websites that served ads, such as how people decided whether or not to interact with an ad, and whether or not they were able to navigate away when they wanted to.

Third, we asked participants to navigate to a blog-style website we created that served six ads taken from our measurement study. As part of our study protocol, we asked participants to refrain from actually clicking on ads, and instead, to talk through their thought process as they decided how they wanted to interact with the ad. Such thought processes included discussing out loud whether they would have decided to click on an ad. These ads included one control ad that we thought was well-designed for accessibility (i.e., had alt-text for images, as well as well-labeled links and buttons), as well as five ads that we hypothesized exhibited at least one inaccessible characteristic (e.g., missing text for links, or containing multiple elements that might "trap" a user's focus). Note that initially, we only included five ads that might have contained inaccessible designs: after interviewing four participants, we decided to place an additional ad on the website. Figures 7-12 show the ads included in our website for the user study, and we summarize each ad's intended "inaccessible characteristic" in its caption. Some ads exhibited more than one inaccessible characteristic.

As participants navigated through our website, we asked them to talk aloud, focusing particularly on the ads they encountered. When appropriate, we prompted with follow-up questions about what cues they had used to determine it was an ad, checking to see if they could understand the content of the ad, and discussing their opinions of its design.

We finished the interview with a wrap-up and reflection section. We gave participants an opportunity to discuss things they felt we had not yet covered. We also asked questions about what advice they would give website owners, advertisement designers, and screen reader companies that might improve the way they browsed online.

Data Analysis. As all of our interviews were conducted over Zoom, the audio was recorded and transcribed through the software. We used an inductive thematic analysis to surface and summarize participants' views on ads. We conducted the interviews, and generated themes based on the quotes and ideas from the participants. We then subsequently collaboratively discussed these organization of the themes, and present the results in the following sections.

Based on prior work and guidelines around qualitative methods, we did not use two coders nor report inter-rater reliability [5, 25]. Importantly, note that we are not making statistical claims about the generalizability of our participants. That is, we are not claiming that all blind participants share the same experiences, or that our participants' experience represent the "average user" of a screen reader. Rather, we use these interviews to surface themes around concerns and preferences among our participants, not quantify how often they encounter ads that they find challenging.



Figure 7: A shoe ad with multiple, unlabeled links



Figure 9: A wine ad with two images that are missing alt-text: a logo, and a turn sign



Figure 11: An carseat ad whose alt-text is non-descriptive (says 'Advertisement')

Figures 7-12 show the ads we included in our website for participants to navigate through during our user study.

Table 7: Participant Demographics

Category	Distribution (Count)
Age	18-24 (6), 25-34 (3), 35-44 (2), 45-54 (1), 55-64 (1)
Gender	Male (7), Female (6)
Race	White (8), Middle Eastern (2), Asian (2), South Asian (1)
Screen reader	NVDA (8), JAWS (6), VoiceOver (11), TalkBack (1)
Years w/ assistive tech	1-5 (2), 6-10 (7), 11-15 (2), 16-20 (2)
Skill level	Advanced (10), Intermediate / Advanced (3)

Participants. We summarize participant's characteristics in Table 7. All of the participants we interviewed spoke English fluently, were fully blind, and relied solely on a screen reader to convey information as they navigated online content. Note that we did not interview people if they used any visual cues to process information (e.g., people who were low-vision, or only had mild vision impairments). On average, participants in our interview were 31 years old, used screen readers for 10 years, and rated themselves as either advanced, or between intermediate and advanced users of the technology. Most of the participants in our study used more than one screen reader, some had multiple installed on their laptops, while others used different screen readers on their laptop and on their smartphone. Though most of our participants were from the



Figure 8: A control, 'well-designed' ad for dog chews



Figure 10: An airline ad with the disclosure in an element that is not keyboard focusable



Figure 12: A bank ad with missing alt for images, and unlabeled buttons

US (12 participants), we also interviewed people from Pakistan (1) and Egypt (2).

6 RESULTS: USER STUDY

We discuss the themes that surfaced from our interviews.

Context: Most people did not use ad blockers. On the whole, our participants browsed the web with ads: of the 13, only three used an ad blocker (two only in the context of work). Most participants cited usability reasons to explain why they did not use an ad blocker. Many found that modern websites can block their content if they detect a user has an ad blocker enabled; disabling the ad blocker in response was cumbersome, and not worth the benefits. Others expressed similar opinions, stating that enabling an ad blocker just meant more steps that they did not want to take.

Context: All participants correctly identified the control ad. All of the participants in our study were able to identify the "well-designed" control ad used in our study as third-party content. They not only observed that it was an ad, but also accurately described its contents and could easily decide how they wanted to interact with it. Though (per our direction) no participants actually clicked on any of the ads on our page, two expressed potential interest, as they owned dogs. The remaining participants commented on how it was straightforward to determine interest, and to navigate away once they decided it was not relevant.

6.0.1 People respond to ads by trying to navigate away, especially if the content is inaccessible. All of the participants in our study shared the opinion that ads distracted, and detracted, from their overall browsing experience. In most cases, they stated that their first reaction, encountering an ad, was to scroll past it as fast as possible. The overwhelming sentiments towards ads were negative: the only positive characteristics people described about ads was when they were easy to close, or navigate away from. Some people said that if an ad was disruptive enough, they would simply close out of the website, and find different resources.

This opinion held even more strongly for inaccessible ads. For example, in the case of missing alt-text, participants were not willing to spend more time trying to figure out what the ad contained. P7 said: “If they (the advertisers) aren’t going to spend time making it easy for me to understand, I’m not going to waste my time. I’m just going to scroll right past.” This is despite the fact that some participants, including P7, knew of techniques that leverage AI to summarize the content of images, such as Be My Eyes [8].

Similarly, participants said they did not spend time trying to find more information if the ad initially provided unclear information. For example, if a link said “click here to learn more,” and included additional information through its title attribute, users would not bother to try to find it. Thus, we find that ads that do not expose information clearly to screen reader users lose potential clicks.

6.1 Understandability

6.1.1 Participants often use context clues to identify ads. In our measurement section, one of the aspects we focused on was determining whether ads used language in text to describe their status as third party content. However, when asking participants, we found that instead of relying on ads to disclose their status through language, everyone instead used context to determine when they were on an ad or interacting with content on a page. While some mentioned that they heard keywords, such as “advertisement,” or “sponsored,” from time to time, it was not the primary way they decided what was an ad, and what was not.

One participant, P8, said “It depends on what I’m expecting. If I’m on a news website, and I suddenly hear something about medicine, I’ll know that the medicine is an ad.” This is why we decided to add an additional “stealthy” ad to our website. We hypothesized that the ad, as it contained a disclosure in an element that did not receive keyboard focus, might be harder for users to know that it was an ad. However, in subsequent interviews, all participants still detected the Alaska Airlines as being an ad. We suspect that this is because there was a mismatch between the content of our website (a blog), and the content of the ad (an airline). Future research could continue this line of questioning: more subtle differences, such as an ad for an airline appearing on a website that displayed travel tickets might be harder to detect for users.

Non-descriptive content confused people. Figure 11 shows the ad on our website whose alt-text simply said “Advertisement,” though it is discussing the importance of choosing correct car seats for children. This represents the behavior shared across 1/3 of the ads in our data set. Every participant in our study was not able to initially detect this ad as being its own ad. Only after we later alerted participants to it did they realize that it was its own ad unit.

This may have been compounded by the fact that the ad is right next to others: many thought that it was part of the ad below it in the sidebar. However, it is common for websites to include multiple ads next to each other.

As with ads that lack alt-text to describe images, we find that ads that only use non-descriptive language are more likely to lose clicks from people who are blind and use screen readers to navigate the web. And, as we highlighted in previous sections, blind users are unable to distinguish the content that might be interesting, from content that uses similar non-descriptive language, but leads to malicious websites.

6.1.2 Unlabeled links confused people. On our test website, Figure 7 represents an ad that is neither understandable nor navigable due to multiple unlabeled links. All participants in our study found this ad to be highly annoying, and none were able to understand what it was trying to promote.

While participants navigated through our site, this ad prompted three (we found) interesting conversations. First, one participant, P12, had their focus “trapped” in the ad, and was unable to actually navigate through to the rest of the blog without using shortcut keys built into the screen reader to jump to the next header element on the page. People who use screen readers rely on well-designed websites: if a page does not have clear landmarks, navigating away from (third-party) focus traps might be impossible. Second, not all users may know about the shortcut keys that would work to navigate them away from such focus traps. In these instances, ads with multiple elements that are unclear to users might deter them from using an otherwise accessible website entirely.

Second, a different participant, P13 was surprised by this ad, because they had initially not thought it was an ad at all. They realized, through participating in the walkthrough of the website, that the behavior of the ad – reading out loud a series of links with numbers and strings – were components of an ad. They mentioned encountering similar behavior in the past, but thought that it was just broken parts of websites that they were trying to visit. This shows that there are instances where people who use screen readers misunderstand whether or not they are interacting with an ad, and that unlabeled links make it even more confusing.

Finally, P4 was able to understand that the ad was delivered by Google, even though they did not understand the content. As they explained, Google ads were so often inaccessible, in the same way (i.e., unlabeled links) that they recognized the pattern of the domain the screen readers announced.

6.2 Navigability

6.2.1 Advertisements that are hard to close or navigate away from frustrated people. All participants reported at least one instance where they felt as though advertisements disrupted the way they browsed online. This included ads that contained too many elements, making it difficult to navigate or scroll past them. Of all the ads presented on our test website, the shoe ad in Figure 7 was the one that people uniformly found most frustrating, because of the number of elements that were in one ad unit. Similarly, ads that were difficult to close – where the x button was not labeled or easily discoverable – were annoying to scroll past.

Participants also discussed problematic ad designs that did not come up directly in our study, such as pop-up ads. Though we did not collect pop-ups in our crawls, many participants found them frustrating because they are difficult to close, and because users do not get their focus back to where they were on the page after closing the ad.

A few participants also described how video ads on cooking websites “yelled” over their screen readers, disrupting their browsing as they scrolled through the page. Instead of hearing their screen reader say the content as they scrolled, they would hear the ad announcing itself repeatedly, counting down the number of seconds until a video ad starts playing, regardless of where participants were on the page. Though we did not observe video ads in our measurement — our collection methods meant that we only saw ads in one snapshot in time — the solution may be straightforward: using ARIA-live polite regions ensures that content cannot override the control of a users’ screen reader.

7 LIMITATIONS

First, for our measurement study, we capture only a particular sample of ads and websites. Our crawling approach relies on EasyList, and uses Chrome to render pages and images. Though there are known limitations to EasyList’s detection of ads, it is commonly used as a method to identify ads in measurement studies [42]. The website categories we selected for crawling was aimed to capture a variety of ads, not to represent a generalizable sample of all ads or websites across the web — there are of course other types of sites with ads (e.g., cooking sites, as mentioned by our user study participants), and future work may wish to compare the accessibility of ads on different types of sites.

As we found that the majority of participants did not use ad blockers, we did not fully explore how ad blockers might help the way people who are blind or have low vision navigate websites. Future work could continue working with participants to understand how using ad blockers changes their ability to access websites and content.

Because we clear cookies and use a clean profile each time we visit a website, the quality of the ads we received may have differed from those seen by users with extensive histories.

We identified ad platforms by relying on visual heuristics, as we did not track or record network requests while loading our pages. This means that we were not able to use network-based methods for identifying which ad platforms deliver ads, such as analyzing inclusion chains outlined by Bashir et al. [7].

We were able to scale our analysis of ads by relying on the ChromeDevTool’s API to access each ad’s accessibility tree. However, we did not crawl with a screen reader enabled, and our experience with screen readers comes from manual testing. Different screen readers can convey different information in different ways. As such, we make broad statements about what screen readers could do, and we note when the behaviors might diverge in meaningful ways.

Finally, we note that our participants skewed young and white, likely because we recruited participants from social media groups, despite extending our invitation to participate to state

groups. After completing our interviews, we heard back from several interested parties in different states, but did not interview more people due to time constraints. We stress again that (as is common in qualitative methods) we do not aim to generalize to the overall population.

8 DISCUSSION

In the interest of equity, we believe that ads should expose enough meaningful information to users, such that everyone understands what it is trying to promote, regardless of how a person navigates the web. This means that people who navigate via keyboard should be able to tell if an ad is relevant to their interests. Simply put, people should not be left out, based on the method they use to browse online. Additionally, our qualitative, semi-structured interviews found that people want to be able to navigate past ads in an easy, straightforward manner that does not disrupt their browsing experience in ways that ads currently do. There are straightforward solutions to both aspects: making sure that ads expose enough information, as well as making ads easier to navigate past.

8.1 Improving Perceivability and Understandability

Ad platforms could create policies that require ads to provide meaningful information to screen readers in the HTML attributes that exist for this purpose (e.g., ARIA-labels, alt-text). We found that many (i.e., more than 75%) of ads that use ARIA-labels or alt-text either leave them blank, or include only generic information. This means that two ads that visually appear quite different would actually seem the same to someone who is using a screen reader. Ad platforms could prevent this behavior by (1) creating a template that encourages the use of assistive attributes, (2) rejecting ads that contain generic strings (or missing attributes), or (2) (potentially) extract more information about the ad even if it is not directly provided by the advertiser. For example, advertising platforms could inspect the meta-property HTML tag of the landing page associated with the ad to provide more information when advertisers provide generic terms.

Advertisers could also make ad content more accessible. They are paying to place their ad in front of audiences they think would be most interested in their product. This includes people who might use screen readers.

Additionally, *websites* should carefully consider the advertising platforms they choose to deliver ads on their page. Our results suggest that some ad platforms tend to serve ads that are more accessible than others.

Indeed, we found that major ad platforms, such as Google, Yahoo, and Criteo, serving inaccessible ads for seemingly straightforward reasons. As such, the solutions are also technically simple. Because the ad ecosystem is largely made up of small number of influential players, making these small changes would have a long-reaching impact.

It is possible that there are other reasons underlying some of these seemingly straightforward accessibility limitations. For example, ads that are more easily programmatically identifiable as ads are also easier for ad blockers to identify and block. Thus, there may be a tension between accessibility to screen readers and to ad blockers.

We urge ad platforms to prioritize accessibility for all users (and note that the inaccessible ads we surfaced in our measurement are already detectable by EasyList, used by many ad blockers).

8.2 Improving Navigability

Our participants all wanted better ways to navigate past ads that did not rely on ad blockers. In this regard, *website owners* could create Bypass Blocks (also known as “skip links”) that allow users to easily skip the content of ads. We found that (though uncommon), some ad campaigns have as many as 40 interactive elements: this means that someone navigating the page would need to tab through the ad content 40 times before accessing the content of the website they visited. Bypass blocks not only provide a way for users to skip ads that are “focus traps,” they also allow users to navigate back out of ads they started interacting with, but no longer desire to.

Screen readers could also help their users skip information more easily. For example, they currently have several shortcuts that allow users to navigate through webpages in a nonlinear fashion. However, we did not observe shortcuts that allowed screen reader users to return to the parent content once inside an iframe. Such a shortcut would provide a way for screen reader users to “back out” of an ad after they start interacting with it. Screen reader companies could also consider building in ad blockers, though this could be technically more challenging than our other suggestions.

9 ADDITIONAL RELATED WORK

There are substantial bodies of work and research communities around both web ad ecosystem measurement and web accessibility; this paper sits at the intersection of these and contributes knowledge to both. For example, on the ad ecosystem side, prior work has studied and measured the prevalence and privacy implications of third-party trackers and advertisers on the web [11, 12, 20, 27, 30, 35], how ads are targeted at individuals [13, 22, 34, 43], and problematic content in ads [4, 39, 41, 42, 44].

On the accessibility side, there have been past measurement studies of accessibility on the web [9, 16, 17, 23], though these focus on analyzing websites, and not advertisements. In 2001, Thompson and Wassmuth [33] analyzed the quality and accessibility of alt-text associated with banner ads found on news websites. Our work seeks to take a broader approach by collecting and analyzing ads from a wider variety of websites, and we examine accessibility characteristics beyond alt-text. More recently, He et al. [18] used a mix-methods approach to analyze the accessibility of advertisements on mobile devices. While our work uses similar mixed-methods approaches to collecting ads and interviewing users, we focus on how ads are displayed through the Document Object Model (DOM) on the web, while He et al. analyze ads displayed either in native elements specific to mobile environments, or through WebView. Moreover, accessibility issues on mobile devices typically focus on users’ ability to swipe or touch, while web accessibility is influenced by keyboard shortcuts and mouse navigation. Work in the human-computer interaction and accessibility communities has also often engaged directly with screen reader users in other contexts [3, 10, 18, 31].

10 CONCLUSION

We ran a 31-day measurement of 90 websites, collecting ads and analyzing their accessibility based on WCAG best practices. We find that significant fractions of ads have inaccessible characteristics. To put the implications of our measurement findings in context, we conducted a qualitative user study, highlighting the ways in which screen reader users currently find ads hard to understand and navigate. Finally, we make suggestions for website owners, advertisers, advertising platforms, and screen readers to make people’s browsing experiences more equitable. We are in the process of reaching out to ad platforms to share our findings, particularly in cases where seemingly simple fixes would have a large positive impact on accessibility to screen readers.

11 ETHICS

Because none of the authors on this paper are blind or use a screen reader, we consulted with accessibility researchers and members of the blind and low vision community at our institution during both the measurement and user study. We did this to ensure that our research questions were grounded in the needs of members of the community, to discuss problematic behaviors people had anecdotally observed in ads and screen readers, and to confirm that the questions we asked participants were reasonable and not overly burdensome.

Measurement Study. Within the measurement subfield focused on online advertising, it is common to crawl websites to either load ads or load and then click on ads [29, 38, 44]. Our measurement study visited and loaded the ads on 90 popular websites once daily over the course of one month. We did not click on any ads. As with prior ad-focused web measurements, we believe the impact on ad impressions from our crawling to be very small compared to the volume of traffic these websites receive daily, and in line with prior work crawling and studying the online ad ecosystem.

Our findings reveal accessibility improvements that can be made to many ads; we have reached out to the major ad platforms (i.e., Google, Criteo, Yahoo). Google is in the process of verifying our findings and updating the accessibility of the ‘Why this Ad?’ buttons. We have reached out to Criteo and Yahoo again, to share our findings and encourage them to improve the accessibility of their ads.

User Study. All user study procedures were approved by our institution’s IRB, and participants gave their informed consent prior to beginning the interview.

12 ACKNOWLEDGEMENTS

We would like to thank our paper shepherd, Waqar Aqeel, and the anonymous reviewers for their insightful comments. We would also like to thank Arnavi Chheda-Kothary, Umar Iqbal, and Ather Sharif for their help brainstorming in the early stages of the project, as well as the participants who made this research possible. This work was supported in part by the US National Science Foundation under Award #2041894.

REFERENCES

- [1] Date accessed: 2024-05-14. *Digital Advertising - US | Statista Market Forecast*. <https://www.statista.com/outlook/dmo/digital-advertising/united-states>
- [2] HeyTony Advertising. 2023. SEO Link Title Attribute Best Practices. <https://heytony.ca/seo-link-title-attribute-best-practices/>
- [3] Khaled Albusays, Stephanie Ludi, and Matt Huenerfauth. 2017. Interviews and Observation of Blind Software Developers at Work to Understand Code Navigation Challenges. In *International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '17)*. Association for Computing Machinery, New York, NY, USA, 91–100.
- [4] Muhammad Ali, Angelica Goetzen, Alan Mislove, Elissa M. Redmiles, and Piotr Sapiezynski. 2023. Problematic Advertising and its Disparate Exposure on Facebook. In *USENIX Security Symposium*.
- [5] David Armstrong, Ann Gosling, John Weinman, and Theresa Marteau. 1997. The Place of Inter-Rater Reliability in Qualitative Research: An Empirical Study. *Sociology* 31, 3 (1997), 597–606. <https://doi.org/10.1177/0038038597031003015> arXiv:<https://doi.org/10.1177/0038038597031003015>
- [6] Eric Bailey. 2017. *The Trials and Tribulations of the Title Attribute*. <https://www.24a11y.com/2017/the-trials-and-tribulations-of-the-title-attribute/>
- [7] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. 2016. Tracing information flows between ad exchanges using retargeted ads. In *25th USENIX Security Symposium (USENIX Security 16)*, 481–496.
- [8] Be My Eyes. Date accessed: 2024-05-06. *Be My Eyes - See the world together*. <https://www.bemyeyes.com/>
- [9] Trevor Bistic, Jeffrey Stanley, John Higgins, Daniel Chudnov, Justin F. Brunelle, and Brittany Tracy. 2021. Automated Evaluation of Web Site Accessibility Using A Dynamic Accessibility Measurement Crawler. *CoRR* abs/2110.14097 (2021). arXiv:2110.14097 <https://arxiv.org/abs/2110.14097>
- [10] Julian Brinkley and Nasseh Tabrizi. 2017. A Desktop Usability Evaluation of the Facebook Mobile Interface using the JAWS Screen Reader with Blind Users. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 61, 1 (2017), 828–832.
- [11] John Cook, Rishab Nithyanand, and Zubair Shafiq. 2020. Inferring Tracker-Advertiser Relationships in the Online Advertising Ecosystem using Header Bidding. In *Privacy Enhancing Technologies Symposium (PETS)*.
- [12] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: A 1-million-site measurement and analysis. In *ACM Conference on Computer and Communications Security (CCS)*.
- [13] Irfan Faizullabhoy and Aleksandra Korolova. 2018. Facebook's Advertising Platform: New Attack Vectors and the Need for Interventions. In *Workshop on Consumer Protection (ConPro)*.
- [14] Federal Trade Commission. 2013. Dot Com Disclosures: Information About Online Advertising. Available at <https://www.ftc.gov/system/files/documents/plain-language/bus41-dot-com-disclosures-information-about-online-advertising.pdf>.
- [15] The Paciello Group. Date accessed: 2024-05-12. *Using the HTML Title Attribute*. <https://www.tpgi.com/using-the-html-title-attribute-updated/>
- [16] Alexander Hambley. 2021. Empirical web accessibility evaluation for blind web users. *SIGACCESS Accessibility and Computing* 129, Article 2 (mar 2021), 5 pages.
- [17] S. Harper and A.Q. Chen. 2012. Web accessibility guidelines: A lesson from the evolving web. *World Wide Web* 15 (2012), 61–88.
- [18] Ziyao He, Syed Fatul Huq, and Sam Malek. 2024. "I tend to view ads almost like a pestilence": On the Accessibility Implications of Mobile Ads for Blind Users. In *IEEE/ACM International Conference on Software Engineering (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 197, 13 pages.
- [19] Amber Hinds. 2023. Divs Are Not Buttons. <https://theadminbar.com/accessibility-weekly/divs-are-not-buttons/>
- [20] Umar Iqbal, Peter Snyder, Shitong Zhu, Benjamin Livshits, Zhiyun Qian, and Zubair Shafiq. 2020. AdGraph: A Graph-Based Approach to Ad and Tracker Blocking. In *In the Proceedings of the IEEE Symposium on Security & Privacy*.
- [21] Satwik Ram Kodandaram, Mohan Sunkara, Sampath Jayarathna, and Vikas Ashok. 2023. Detecting Deceptive Dark-Pattern Web Advertisements for Blind Screen-Reader Users. *Journal of Imaging* 9, 11 (2023), 239.
- [22] Mathias Lécuyer, Riley Spahn, Yannis Spiliopolous, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. 2015. Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidence. In *ACM Conference on Computer and Communications Security (CCS)*.
- [23] Rui Lopes, Daniel Gomes, and Luís Carriço. 2010. Web not for all: a large scale study of web accessibility. In *International Cross Disciplinary Conference on Web Accessibility (W4A)* (Raleigh, North Carolina) (*W4A '10*). Association for Computing Machinery, Article 10, 4 pages.
- [24] John Mahoney. 2015. A Complete Taxonomy of Internet Chum. *The Awl*. <https://www.theawl.com/2015/06/a-complete-taxonomy-of-internet-chum/>.
- [25] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (nov 2019), 23 pages. <https://doi.org/10.1145/3359174>
- [26] Jeremy B. Merrill. 2020. Facebook didn't mark ads as ads for blind people for almost 2 years. <https://qz.com/1800064/for-blind-facebook-users-ads-havent-been-labeled-as-ads>
- [27] Nick Nikiforakis, Alexandros Kapravelos, Wouter Joosen, Christopher Krügel, Frank Piessens, and Giovanni Vigna. 2013. Cookieless Monster: Exploring the Ecosystem of Web-Based Device Fingerprinting. *IEEE Symposium on Security and Privacy* (2013), 541–555.
- [28] The A11Y Project. 2013-04-22. *Title Attributes*. <https://www.a11yproject.com/posts/title-attributes/>
- [29] Vaibhav Rastogi, Rui Shao, Yan Chen, Xiang Pan, Shihong Zou, and Ryan Riley. 2016. Are these Ads Safe: Detecting Hidden Attacks through the Mobile App-Web Interfaces. In *NDSS*.
- [30] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. 2012. Detecting and Defending Against Third-Party Tracking on the Web. In *USENIX Symposium on Networked Systems Design and Implementation (NDI)*.
- [31] Ather Sharif, Sanjana Shivan Chintalapati, Jacob O. Wobbrock, and Katharina Reinecke. 2021. Understanding Screen-Reader Users' Experiences with Online Data Visualizations. In *International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21)*. Association for Computing Machinery, New York, NY, USA, Article 14, 16 pages.
- [32] The World Wide Web Consortium. Date accessed: 2024-05-05. W3C. <https://www.w3.org/>
- [33] David Thompson and Birgit Wassmuth. 2001. Accessibility of online advertising: a content analysis of alternative text for banner ad images in online newspapers. *Disability Studies Quarterly* 21, 2 (2001).
- [34] Giridhari Venkatadri, Athanasios Andreou, Yabing Liu, Alan Mislove, Krishna P. Gummadi, Patrick Loiseau, and Oana Goga. 2018. Privacy Risks with Facebook's PII-Based Targeting: Auditing a Data Broker's Advertising Interface. In *IEEE Symposium on Security and Privacy*.
- [35] Paul Vines, Franziska Roesner, and Tadayoshi Kohno. 2017. Exploring ADINT: Using Ad Targeting for Surveillance on a Budget - or - How Alice Can Buy Ads to Track Bob. In *Workshop on Privacy in the Electronic Society (WPES)*.
- [36] Similar Web. Date accessed: 2024-01-06. Effortlessly Analyze Your Competitive Landscape. <https://www.similarweb.com/>
- [37] WHATWG. Date accessed: 2024-02-12. *HTML Living Standard: The title attribute*. <https://html.spec.whatwg.org/multipage/dom.html#the-title-attribute>
- [38] Xinyu Xing, Wei Meng, Byoungyoung Lee, Udi Weinsberg, Anmol Sheth, Roberto Perdisci, and Wenke Lee. 2015. Understanding Malvertising Through Ad-Injecting Browser Extensions. In *24th International Conference on World Wide Web (WWW)*.
- [39] Christina Yeung, Umar Iqbal, Yekaterina Tsipenyuk O'Neil, Tadayoshi Kohno, and Franziska Roesner. 2023. Online Advertising in Ukraine and Russia During the 2022 Russian Invasion. In *The Web Conference (WebConf)*.
- [40] Eric Zeng. Date accessed: 2023-12-10. AdScraper. <https://github.com/UWCSESecurityLab/adscraper>
- [41] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. 2020. Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites. In *Workshop on Technology and Consumer Protection (ConPro)*.
- [42] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. 2021. What makes a "bad" ad? user perceptions of problematic online advertising. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–24.
- [43] Eric Zeng, Rachel McAmis, Tadayoshi Kohno, and Franziska Roesner. 2022. Factors Affect Targeting and Bids in Online Advertising? A Field Measurement Study. In *ACM Internet Measurement Conference (IMC)*.
- [44] Eric Zeng, Miranda Wei, Theo Gregersen, Tadayoshi Kohno, and Franziska Roesner. 2021. Polls, Clickbait, and Commemorative \$2 Bills: Problematic Political Advertising on News and Media Websites Around the 2020 U.S. Elections. In *ACM Internet Measurement Conference (IMC)*.

A USER STUDY PROTOCOL

Our full interview protocol included the following questions. Note that since this was a semi-structured interview, the interviewer may have asked slightly altered or additional questions depending on the particular conversation.

Background.

- (1) What platform do you do most of your web browsing (Desktop, Laptop, Phone)?
- (2) Which browser + OS do you use?
- (3) What types of assistive technologies do you use when browsing online services? What are the names of the tools you use?
- (4) Why do you use those assistive technologies? How do these technologies help you as you navigate, compared to how you would browse the web without them?
- (5) How long would you say you've been using [insert name] assistive technology?
- (6) Would you rate your expertise with [insert name] assistive technology as Novice, Intermediate or Advanced?
- (7) How many hours of online browsing do you do each day (on average)? [None at all, More than 0 but less than 1 hour a day, more than 1, but less than 3, more than 3 but less than 5, more than 5]
- (8) What types of online services do you commonly use (e.g., shopping sites, airlines, online banking, news, etc.)?

Experience with ads.

- (1) Have you heard about ad blockers? Do you use an adblocker when navigating content online? If yes: Why? If no: Why not?
- (2) What type of ads do you typically come across during browsing?
- (3) Can you talk a bit about your experiences encountering ads as you navigate websites?
- (4) Is there anything that annoys you about any ads you've encountered, or things that you've liked?
- (5) What is your initial reaction when you encounter an ad?
- (6) Are there specific cues you use to identify when you're interacting with an ad, instead of the content of the page you intended to visit?

- (7) Does it make a difference if ad disclosures are in elements that are not keyboard focusable? When ad disclosures appear later in the ad, do you feel like it's misleading?
- (8) How often do you choose to click on ads? Do you ever click on ads accidentally?
- (9) How do you decide whether it's safe or not to click on an ad?
- (10) When interacting with something you know is an ad, do you think the ad provides sufficient details such that you know what it's conveying?
- (11) How often do you choose to engage with descriptions, when they're available? When you do interact with descriptions, do you find that it contains useful information?
- (12) How much do you rely on alt-text? What do you do if there is no alt-text? How often do you feel as though you are not receiving information you need in order to make decisions about interacting with the content?
- (13) Are there other strategies you use, like asking Google to identify what is in the image?
- (14) Have you encountered ads that have too many elements, or "trap" your focus? If so, how do you navigate away from such ads?
- (15) Does the location of an ad on a web page affect your ability to detect an ad, or interact with it?

Interacting with our website (see Figures 7-12). For this part, I'd like you to visit a page we've created: there will be some things that we've designed to mimic real-world "ads" that we've observed. I'd like you to navigate the page, and just say what you're thinking out loud as you're browsing through.

Reflection and wrap-up.

- (1) Is there anything you would like website designers, or online ad designers, or the designers of accessibility tools to know about your experience with ads as a screen reader user?
- (2) Have you felt as though ads affect your ability to browse websites? If so, how? If not, why not?
- (3) (If they use JAWS) Did you know that there's a built-in feature in JAWS that allows you to skip content in iframes (which typically contain ads)? If yes: Do you enable this feature? If no: Would you want to enable this feature / does it sound like something that would make web navigation easier?
- (4) Is there anything else you'd like to share with me?