The Value of First Impressions
The Impact of Ad-Blocking on Web QoE

James Newman Fabián E. Bustamante

Northwestern University

Abstract. We present the first detailed analysis of ad-blocking's impact on user Web quality of experience (QoE). We use the most popular web-based ad-blocker to capture the impact of ad-blocking on QoE for the top Alexa 5,000 websites. We find that ad-blocking reduces the number of objects loaded by 15% in the median case, and that this reduction translates into a 12.5% improvement on page load time (PLT) and a slight worsening of time to first paint (TTFP) of 6.54%. We show the complex relationship between ad-blocking and quality of experience - despite the clear improvements to PLT in the average case, for the bottom 10 percentile, this improvement comes at the cost of a slowdown on the initial responsiveness of websites, with a 19% increase to TTFP. To understand the relative importance of this trade-off on user experience, we run a large, crowd-sourced experiment with 1,000 users in Amazon Turk. For this experiment, users were presented with websites for which ad-blocking results in both, a reduction of PLT and a significant increase in TTFP. We find, surprisingly, 71.5% of the time users show a clear preference for faster first paint over faster page load times, hinting at the importance of first impressions on web QoE.

1 Introduction

The web advertisement industry has grown exponentially over the past decade and is now the primary source of income for most content providers [19]. A number of research efforts in the last few years have focused on understanding their scale, mechanisms, and economics [10, 26, 34].

While keeping most online content and services "free", web advertisements have raised serious security problems and privacy concerns, and attracted some negative press due to questionable practices [21, 29, 40]. In response, millions of users have adopted some form of ad-blocker. By February 2017, at least 615 million devices have an ad-blocker installed, and the total ad-block usage increased 30% between December 2015 and 2016, according to the latest PageFair Adblock Report from 2017 [8].

Besides increased security and fewer interruptions, a key motivation for the wide adoption of ad-blockers is speed [8]. While it seems intuitive that loading fewer objects would lead to an improved quality of experience (QoE), the exact impact of ad-blocking on perceived website performance is still unclear, despite the importance of QoE on user engagement and profit [24, 39].

This paper presents the first detailed analysis of ad-blocking's impact on user QoE. We use the most popular web-based ad-blocker to analyze the impact of ad-blocking on users' web experience when visiting the top-5,000 most popular websites according to Alexa. We rely on three commonly used metrics as proxies of users' QoE – Page Load Time, First Paint Time, and Speed Index.

Our results reveal a complex relationship between ad-blocking and web QoE. We find, as expected, that ad-blocking reduces the number of objects loaded by 15% in the median case. This reduction in loaded objects translates into a 12.5% improvement on PLT and a slight worsening of TTFP of 6.54%, on average. When focusing on the bottom 10 percentile, however, we find that while ad-blocking yields a 14% improvement on PLT, the worsening of TTFP is about 3x higher than in the average case.

To understand the relative importance of this trade-off for user experience, we conducted a large crowd-sourced experiment of ad-blocking and Web QoE with 1,000 users in Amazon Mechanical Turk. Users were presented with websites for which ad-blocking results in both a significant reduction of PLT and a significant increase of TTFP. We find, surprisingly, that 71.5% of the time users show a clear preference for faster first paint over page load times, suggesting the importance of first impressions on web QoE.

In summary, our main contributions are:

- We report on the first large-scale evaluation of the web QoE impact of ad-blocking with the 5,000 top Alexa sites.
- We show the complex relationship between ad-blocking and quality of experience – while ad-blocking yields clear improvements to PLT in the average case, for the bottom 10 percentile, this improvement comes at the cost of a significant slowdown on the initial responsiveness of websites, with a 19% increase to TTFP.
- We present results from the largest crowd-sourced analysis of ad-blocking impact on QoE today, with 1,000 users in Amazon Mechanical Turk. Our results suggest that user experience is more sensitive to faster first paint than slower page load times.
- To assist open science, we will publicly release our dataset from our controlled experiment with the top 5,000 Alexa websites and the 1,000-user crowd-sourced experiment.[1]

2 Ad-Block Background

Ad-blockers come in a number of formats – as browser extensions, VPN-based solutions and full browsers (e.g., Brave, Cliqz and now Chrome Canary [13]). In this work we focus on the browser-extension format as this is by far the most commonly used option.

There is a wide range of browser extensions aimed at avoiding or blocking ads including Ghostery, 1Blocker, NoScript, Adblock, and Adblock Plus [18].

[1] `http://www.aqualab.cs.northwestern.edu/projects/AdQoE`

Most web-based ad-blockers rely on the browser's webRequest API to intercept requests from websites for modification [7]. The API allows an extension or plugin to act as a proxy and interact with requests from the website at different points in their life cycle.[2] For example, in the Chrome browser Adblock Plus uses the "onBeforeRequest" callback to receive the URL of a request and determines whether or not to block it.

To decide whether a URL should be blocked or not, ad-blockers use crowd-sourced list of filter rules ("filter lists"). Filter list rules are regular expressions that match HTTP requests and HTML elements. Ad-blockers block HTTP requests and hide HTML elements if they match any of the filter rules. Ad-blockers typically allow users to subscribe to different filter lists and even incorporate custom filter rules. EasyList [32] is the most popular of these list, but there are others such as Fanboys Enhanced Tracking List [11], Disconnect.me [9] and Blockzilla [14] as well as language-specific ones [3]. Filters lists can include thousands of rules; EasyList alone is over 69K-rules long at the time of submission.

For our analysis we use Adblock Plus (ABP). ABP is by far the most popular ad-blocker holding, according to a recent study by Malloy et al. [19], over 90% of the market for Firefox and Internet Explorer and nearly 50% of Google Chrome's market. Despite or focus on ABP, we believe our finding are generalizable to any of the ad-blockers relying on EasyList or similar filter lists.

3 The Performance Cost of Ad-Blocking

Our analysis aims to identify specifically how ad-blockers impact user QoE. The following paragraphs present our experimental methodology and dataset, and describe our evaluation results.

3.1 Methodology and Dataset

To analyze the impact of ad-blocking on user QoE, we load a range of popular websites in a controlled environment, with and without ABP enabled. For this we use WebPageTest (WPT) [20], an online, open-source web performance diagnostic tool. WPT creates a sandbox with virtual machines in which testers can load websites using various devices and browsers over different network conditions. The tool returns a straightforward report card summarizing the performance results of its tests, including a table of milestones alongside speed metrics, such as PLT and TTFP. WPT performs a similar analysis of web pages as [30] and was used by Netravali et al. for accurate record-and-replay for HTTP [23].

We employ a private instance of WPT using a dedicated virtual machine on a desktop and a web server instantiated on Google Cloud Platform [1]. We use a private instance instead of a public one to avoid polluting our results

[2] https://developer.chrome.com/extensions/webRequest

with traffic from other users' concurrent tests. To load different websites in succession, with and without Adblock Plus enabled, we use the Chrome browser flag `--load-extension` to load the unpacked Adblock Plus extension from the local computer's storage.

For our analysis, we aim to limit overhead or bias from the testbed. As such, we do not add latency and leave the default bandwidth for both upload and download. We use the Google Chrome browser (version 57.0.2979.2[3]) with Adblock Plus browser extension (version 1.12.4). After all of the websites have completely loaded, we collect the results using the WPT REST API hosted on our web server, before parsing the resulting HAR files with Haralyzer [12].

We use the top 5,000 popular sites world-wide according to Alexa [2]. This set includes websites with similar URLs and different country codes, which we opted to keep as the may be hosted by different servers or CDNs and potentially be affected differently by ad-blocking.

3.2 Ad-Blocking, Requested Objects and Web QoE

We begin by measuring the reduction in the number of objects loaded with ad-blocking for the top 5,000 Alexa sites. We focus on two of the metrics, Page Load Time and Time to First Paint.[4]

Impact on Requested Objects Sites are made up of many different types of objects including HTML, CSS, JavaScript, image, and video files. Butkiewicz et al. [5] highlight the growing complexity of websites and report that loading a base web page requires fetching more than 40 objects in the median case. For a non-trivial fraction (20%) of websites, the number of objects requested is well above 100. Our analysis on the impact of ad-blocker usage on the number of requested objects focuses on three of the five common object types – JavaScript, images, and HTML [22] – as these are most typically associated with ads [25].

When websites are loaded with ABP enabled, we see 19% fewer objects requested on average, and a 75% reduction for the 95pct (220 instead of 900 objects when loading with ABP disabled).

Table 1 shows a set of percentile numbers of requested objects across our collection of web sites, when the sites were loaded with and without ad-blocker. The last column in the table is the ad-block exposure rate, defined as in Malloy et al. [19], as the number of ads shown to ad-block users per ad shown to no-ad-block users. The drop in the number of requested objects is clear; at the 30th percentile there is already a 10% reduction from ad-blocking. At the 90th percentile, the use of ad-blocking yields a 25% reduction in ad-block exposure rate.

In Figure 1 we focus on the difference in requests for various object types across the 90th percentile of websites. The most blocked type of objects are

[3] The newest version able to work with WebPageTest.

[4] We excluded SpeedIndex results for space considerations; these results were consistent with other findings.

Percentile	Adblock	No Adblock	Ad-Block Exposure Rate
10	20.0	20.2	0.99
30	51.0	56.0	0.91
50	79.0	93.0	0.85
70	116.0	140.0	0.83
90	177.0	237.0	0.75

Table 1: When loaded with ad-block, websites request noticeably fewer objects. This is captured by the ad-block exposure rate, the number of ads shown to an ad-block user per ad shown to a non-blocking user.

images, followed by JavaScript and HTML files. We see that, on average, 5 fewer requests were made for images, while 3 fewer JS and 2 fewer requests for HTML objects were made. This is, in many ways, as expected since images make up the core of ads [25], and are typically requested asynchronously by JavaScript or included in HTML pages.

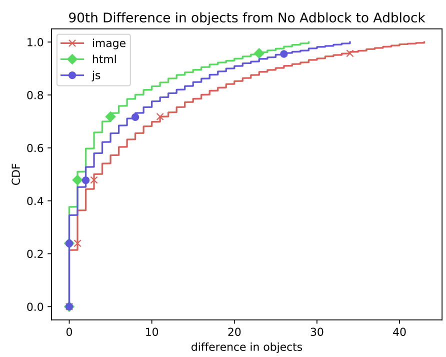

Fig. 1: 90th percentile of the number of requests for types of objects most commonly associated with ads. The most blocked type of objects are images, followed by JavaScript and HTML objects.

Overall Impact on QoE In the following paragraphs we focus on the impact that the decreased number of objects requested has on user QoE as captured by TTFP, PLT, and SI.

Page Load Time (PLT). Page load time (PLT), the most ubiquitous QoE metric, is an approximation of the time it takes for all objects on the website to load. PLT is typically measured as the time between when a page is requested and when the *OnLoad* event is fired by the browser. While some studies have explored other estimates of PLT (such as perceived PLT [15]), we use the traditional PLT metric as our proxy for user QoE.

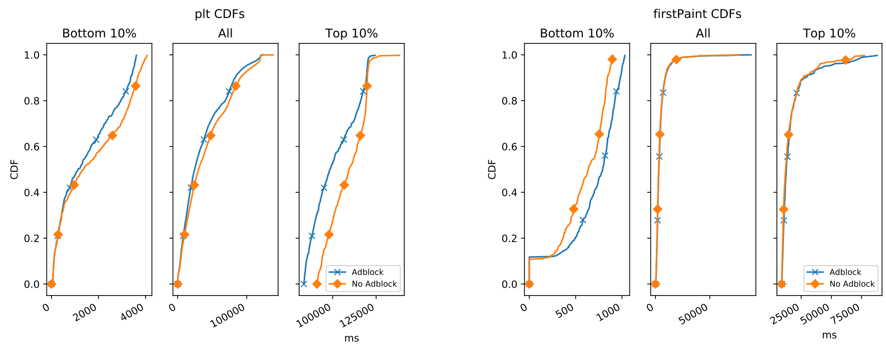

(a) Page Load Time (b) First Paint

Fig. 2: CDFs of Page Load Time (2a) and First Paint (2b; each set includes the bottom 10% and top 10%.

Figure 2a shows the CDFs of PLT with the complete distribution in the center and the bottom and top 10% of the distribution at the left and right of the figure, respectively. The average PLT is approximately 40 seconds when loaded without ABP. Using ABP yields and improvement of 5 seconds, in average. The improvement is even more noticeable in the bottom ten percentile (left-most graph of Figure 2a), with websites loading 14% faster with ABP enabled, as a consequence of loading fewer objects.

Time to First Paint (TTFP). The impact of ad-blocker usage on Time to First Paint is quite different. TTFP captures the time it takes to begin rendering the first objects of a website [35]. When a user navigates to a website, the browser requests the initial HTML page before requesting and rendering the content. TTFP is a function of the complexity of the webpage and the latency to servers hosting the content and is considered an important factor of web QoE, as a lower TTFP means less time a user must wait before starting to view an active site.

Figure 2b shows the CDF of TTFP in milliseconds. The figure contains a similar set of three graphs as with PLT, with the whole distribution in the center and the bottom and top 10% of the distribution at each side of the figure.

The figure shows the clearly negative impact of ad-blocker usage on first paint time, particularly for the lower 10% of the distribution. This delay results from the time it takes for the ad-blocker to decide whether or not it should block an object. Even if the absolute time to process URLs through the EasyList is small, this small overhead can have a significant impact on TTFP for the fastest sites, many of which finish the painting of the first object in less than a second without ad-blocking.

Summary. The use of ad-blocker introduces a constant processing overhead from checking each URL request with the filter list. For many sites, the reduction in the number of ads' associated objects requested yields clear improvements on

PLT. As most ads are loaded asynchronously with JavaScript, however, these benefits do not offset the processing overhead by the time of painting of the first object (TTFP). The following paragraphs explore this trade-off.

4 Crowd-sourced Evaluation of Trade-offs in QoE

The results from the previous section show a clear trade-off in the use of ad-blocker between the *responsiveness* of a website and the total time the user spend waiting on a page to load – for a large number of sites, ad-blocking improves PLT at a significant cost on TTFP.

The relative importance of these two metrics to overall users' QoE, however, is not well understood. To explore this we run a large, crowd-sourced experiment of Web QoE; the following paragraphs describe our experimental methodology and present a summary of our findings.

4.1 Crowd-sourced Experiment Methodology

Our goal is to capture the impact of the trade-off between PLT and TTFP on users' perception. To this end, we need the ability to present a large random set of users with both version of a website, with and without ad-blocker, under the same or similar network conditions.

Experiment Setting. We conduct a user study with 1,000 users on Amazon Mechanical Turk.[5] In our experiments, we direct workers to a website under our control and present them with two versions (with and without Adblock Plus) for each of a sample of sites.[6] For each Human Intelligence Task (HIT), a user is presented with both versions of 10 sites, loaded with and without ad-blocking, and asked to select the site that "loaded faster."

The websites we use in these experiments were selected, at random, from a subset of 965 websites, from our corpus of 5,000, that show both a significantly slower TTFP and a faster PLT when loaded with Adblock Plus.

As in Varvello et al. [33], rather than using live sites during these experiments, we collect videos of the websites loading through WebPageTest under controlled conditions. Videos recorded with WebPageTest have the time included and end with a gray tinted frame. We modified the server to remove these and make sure the user has no indication of when the website has loaded. We use these videos to provide a consistent experience to all participants, regardless of their network connections and device configurations.

We use two different types of instructions during an experiment to ensure we capture the proper response. The first set of instructions, or *primer*, informs the user as to what they should be looking for during the experiment, asking them to *Immediately select the video they believe loaded first.*

[5] https://www.mturk.com
[6] adblock.aqualab.cs.northwestern.edu

The second set of instructions, or *directions*, instructs users on how to make their actual selection by asking them: *Once one of the websites finishes loading, immediately click the video.*

Pre- and Post-Experiment Survey. Before each experiment, we collect some basic demographic information on users' including gender, age group and country of residence. In addition, we ask two additional questions regarding their familiarity with technology: the range of hours spent online on a typical day, and their own rating of their personal technological expertise. These last questions look to determine the impact a user's perceived level of technical proficiency and experience has on their sensitivity to the performance changes introduced by ad-block.

We also include an exit survey that users must complete before submitting their HIT to Amazon. We ask users whether, in their selection of the page which loaded first, they opted for the page which first showed content or the page which appeared to have loaded everything first. Here we are interested in determining what effect, if any, the user's interpretation of *loaded faster* has on their selection.

Quality Control. We apply a number of common techniques to validate the quality of our crowd-sourced data. First, we restrict our survey to workers that have completed ≥ 50 HITs and have an approval rating $\geq 95\%$. Second, beyond our 10 sites, sampled from a larger set of websites where ad-blocking impact on QoE is ambiguous, we include 2 other websites as control cases. Both cases, placed randomly among the other 10 sites, present an obvious choice of "loading faster" in either the right or left of the screen. We employ this as a form of quality control on all of the HITs. All 1,000 HITs correctly chose the control cases. We received 1,080 experiment results and eliminated 5.6% (52) of them that were partially completed (rated less than 10 sites).

Ethical Considerations. Amazon's conditions of use explicitly prohibit tasks that gather personally identifiable information (PII). The information we did collect is coarse enough that we have no reasonable way to map it to individuals. Our experiments collect data "about what", rather than "about whom", through the relatively innocuous task of selecting videos. Our institution's Institutional Review Board (IRB) did not consider our experiment human subject research.

4.2 Summary of Results

Looking at the set of individual experiments, summarized in Table 2, we see that users chose the website loaded without ad-blocker 71.5% of the time as the one "loading faster". Focusing on users rather than individual page comparisons, we find that 86.7% of them choose the non-ad-blocker option as "loading faster" the majority of the time.

This clear preference for non-ad-blocking appears to be independent of any user attribute, including their age, gender, locale and even their self-reported technical proficiency. As an example, Figure 3a shows users' majority preference,

Indicator	Non-ad-blocking (%)	Ad-blocking (%)
Experiments (10,000)	71.5	28.5
Users (1,000)	86.7	13.3

Table 2: Experiments and user majority preference for ad-blocking and non-ad-blocking. In 71.5% of tests users selected non-ad-blocking as loaded faster.

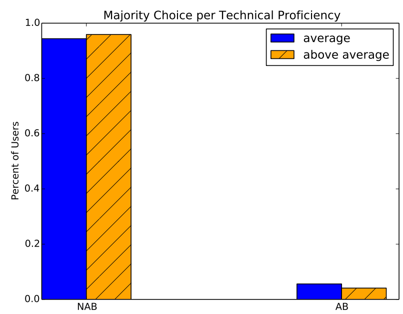

(a) Majority Choice by Technical Proficiency

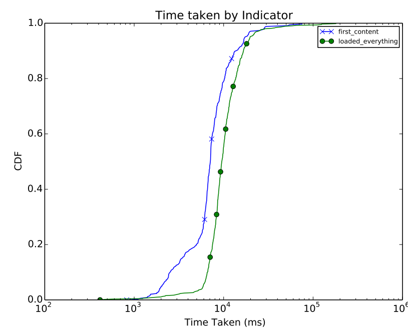

(b) Time Taken

Fig. 3: On the left is the number of users per majority choice aggregated by technical proficiency. On the right is the time taken per user for users based on what they looked for to indicate the page was done loading.

broken down by technical proficiency. The figure combines users with *below average* and *average* self-reported technical proficiency, as only 3.8% of users selected the former. When examining the results of users' preferences by technical proficiency, there is not significant difference in their preferences, with "average" proficiency and "above average" proficiency choosing the non-ad-blocking option 95% and 93% of the time, respectively.

Interestingly enough, this clear preference is even independent of the user's own understanding of what they consider "loading faster" – the site that show some content first or the one that loaded everything first. When aggregating users based on this, we find that despite 65.1% of users selecting *loaded_everything* as their definition of "loading faster", a large fraction of them opted for the non-ad-blocking option that yields a shorter time to first-content (and longer loading time). Table 3 shows the number of users that selected majority non-ad-blocking ("NAB") or ad-blocking ("AB"), separated by their definition of "loading faster" from the exit survey. Over 96% of users selecting "loading everything" opted for the non-ad-blocking (NAB) version.

For validation, Figure 3b plots the time users take to make a decision, aggregated by their choice of "loading faster" in the exit survey. As expected, users who are looking for the first content take less time to select a webpage than users who are looking for everything to be loaded.

Indicator	NAB Users	AB Users	No Majority	Skipped Survey
first_content	227	15	209	17
loaded_everything	525	21		

Table 3: Number of users for "NAB" and "AB" majority choice based on their indicator of page load. We see that users selected *loaded_everything* more often than *first_content*.

5 Discussion

Our measurement study motivates and provides direction to work on improving the quality of experience for users of ad-blockers. The observations are not restricted to the particular ad-blocker we used, but are equally relevant to other ad-blocker that relies on filter lists. Any of these ad-blockers would need to check the list of regular expressions to determine whether or not to block a requested object. As we show, while this check may result in fewer objects being requested and, thus, lower PLT, the extra time will negatively impact TTFP.

There is a number of possible paths to optimize ad-blockers based on their impact on QoE. One could imagine using historical data to identify when loading a website with ad-blocking results in a significantly degraded TTFP / improved PLT. For these sites, the ad-blocker could delay checking until after reaching the TTFP not to impact a website's initial responsiveness if a potential cost on ad exposure.

Our analysis of ad-blocking's impact on users' web experience is preliminary. Our study focuses on how ad-blocking impacts QoE performance metrics such as PLT and TTFP, but that is only one aspect of the whole web browsing experience. Issues such as *Do ad-blockers make web browsing less distracting for users?*, *Do ad-blockers improve users' data privacy?*, or *Despite their performance overhead, do users prefer browsing with ad-blocking enabled?* are interesting research questions that we leave for future work.

6 Related Work

The rapid proliferation of tools to evade or block ads and their potential impact on the web ecosystem have served as motivation to a number of recent studies.

Pujol et al. [26] examines network-wide advertisement traffic and infers the prevalence of ad-block usage. The authors identify advertisement traffic from passive network measurements in a residential broadband network of a European ISP in order to assess the prevalence of ad-blockers. They found that 18% of the total requests in the traffic they monitored were ad-related traffic.

Malloy et al. [19] studies the global prevalence and impact of ad-blockers. Utilizing a dataset composed of information from 2 million users and more than 20 billion page views across half a million top level domains, the paper examines the pervasiveness of ad-blockers around the world. In addition to studying the geographic, demographic, and publisher trends of ad-blockers, the financial impact of ad-blockers on a small set of publishers is discussed. The paper finds that ad-blockers can significantly impact the revenue of publishers, causing a \$3.9M/mo. negative impact on a particular publisher.

A recent study by Walls et al. [34] focuses on Adblock Plus and the Adblock Plus' Acceptable Ads program. This program allows some advertisement providers to pay in order to have their advertisements shown to users. The authors measure the effects of this "whitelist" in order to understand who benefits from it as well as how users perceive "acceptable" advertisements. After running a user study to see how users perceive advertisements, they find that not all advertisements in Adblock Plus' Acceptable Ads program abide by the program's stated policies.

Additionally, different works have analyzed how filter lists work, particularly with respect to anti-adblocking [37]. These works focus on understanding how filter lists identify anti-adblocking functionality on websites. However, they don't examine the time it takes to process the regular expressions present in filter lists.

Other recent work has explored how to effectively defend against JavaScript-based advertisements [10, 16]. These studies attempt to define ways to block JavaScript ads without compromising the security of webpages which serve them. They find that a small number of rules is capable of blocking a large majority of ads on the web.

Our work focuses on a so-far ignored potential side effect of ad-blocker usage: their impact on users' QoE. Internet QoE, and web QoE in particular, have received significant attention in recent years. Much of the work has focused on improving Web page loads, with new network protocols [17, 31], new Web architectures [22, 27, 28, 36], and developing tools. More recent work, such as Kelton et al. [15] and Butkiewicz et al. [6], present alternative approaches and non-traditional metrics to model users' QoE of experience.

7 Conclusions

The growing prevalence of online advertisements has motivated a number of research efforts to understand their scale, mechanisms, and economics while, concurrently, fueling the adoption of services to block them. We presented the first detailed analysis of ad-blocking's impact on user Web QoE. We used the most popular web-based ad-blocker to capture the impact of ad-blocking on common metrics of QoE for the top Alexa 5,000 websites

We found that, while ad-blocking reduction on the number of objects loaded yields a clear improvement on page load time (PLT), for a significant fraction of sites this PLT improvement comes at a high cost on time to first paint (TTFP). We presented results from a large crowdsourced experiment with 1,000 AMT users to understand the relative importance of these metrics on users' experience. We found that, surprisingly, 71.5% of times users indicated a preference for faster TTFP over shorter PLT, hinting at the importance of first impressions on web quality of experience.

While extensive, our evaluation focused on just one aspect of web quality of experience. The impact of ads or the costs/benefits of ad-blockers are not restricted to the chosen metrics of experience we used in our analysis. Understanding other aspects of experience with ad-blocking is left as future work. We have also started to explore ways to leverage our findings to optimize ad-block users' experience, something that will become increasingly relevant as browsers begin to move towards blocking ads [4, 38].

References

1. Google cloud computing, hosting services & apis — google cloud. https://cloud.google.com/.
2. Alexa. The top 500 sites on the web. `https://www.alexa.com/topsites`, 2018.
3. Barrett, C. Filter lists. `https://filterlists.com/`, 2018.
4. Brave Software Inc. Brave browser: Secure, fast and private web browser with adblocker. `https://brave.com/`.
5. Butkiewicz, M., Madhyastha, H. V., and Sekar, V. Understanding website complexity: Measurements, metrics and implications. In *Proc. of IMC* (2011).
6. Butkiewicz, M., Wang, D., Wu, Z., Madhyastha, H. V., and Sekar, V. KLOTSKI: reprioritizing web content to improve user experience on mobile devices. In *Proc. of USENIX NSDI* (2015).
7. Chrome. Chrome webrequest api. `https://developer.chrome.com/extensions/webRequest`. Accessed on 05.25.2018.
8. Cortland, M. 2017 Adblock Report. `https://pagefair.com/blog/2017/adblockreport/`, 2017.
9. Disconnect. Disconnect.me. `https://disconnect.me/`.
10. Dong, X., Tran, M., Liang, Z., and Jiang, X. Adsentry: comprehensive and flexible confinement of javascript-based advertisements. In *Proc. of ACSAC* (2011).
11. Fanboys. Fanboys enhanced tracking list. `https://fanboy.co.nz/filters.html`. Accessed on 05.24.2018.
12. Foundation, P. S. haralyzer 1.4.11. `https://pypi.org/project/haralyzer/`, 2017.
13. Google. Chrome canary. `https://www.google.com/chrome/browser/canary.html`. Accessed on 05.25.2018.
14. Jimdo. Blockzilla: Ad blocking list. `https://blockzilla.jimdo.com/`, 2018.
15. Kelton, C., Ryoo, J., Balasubramanian, A., and Das, S. R. Improving user perceived page load times using gaze. In *Proc. of USENIX NSDI* (2017).
16. Krammer, V. An effective defense against intrusive web advertising. In *Proc. of Conf. on Privacy, Security and Trust* (2008).
17. Langley, A., Riddoch, A., Wilk, A., Antonio Vicente, C. K., Zhang, D., Yang, F., Kouranov, F., Ian Swett, J. I., Bailey, J., Dorfman, J., Roskind, J., Jo Kulik, P. W., Tenneti, R., Shade, R., Ryan Hamilton, V. V., Chang, W.-T., and Shi, Z. The QUIC transport protocol: Design and interent-scale deployment. In *Proc. of ACM SIGCOMM* (2017).
18. Lifehacker. Ad blocking. `https://lifehacker.com/tag/ad-blocking`. Accessed on 05.25.2018.
19. Malloy, M., Matthew, M., Cahn, A., and Barford, P. Ad blockers: Global prevalence and impact. In *Proc. of IMC* (2016).
20. Meenan, P. Webpagetest. `http://www.webpagetest.org/`, 2018. Accessed on 05.24.2018.

21. Metwalley, H., Traverso, S., Mellia, M., Miskovic, S., and Baldi, M. The online tracking horde: a view from passive measurements. In *Proc. of TMA* (2015).
22. Netravali, R., Goyal, A., Mickens, J., and Balakrishnan, H. Polaris: Faster page loads using fine-grained dependency tracking. In *Proc. of USENIX NSDI* (2016).
23. Netravali, R., Sivaraman, A., Das, S., Goyal, A., Winstein, K., Mickens, J., and Balakrishnan, H. Mahimahi: Accurate record-and-replay for HTTP.
24. Poss, T. How does load speed affect conversion rate? `https://blogs.oracle.com/marketingcloud/how-does-load-speed-affect-conversion-rate`, January 14, 2016.
25. Post, E. L., and Sekharan, C. N. Comparative study and evaluation of online ad-blockers. In *Proc. of International Conference on Information Science and Security* (2015).
26. Pujol, E., Hohlfeld, O., and Feldmann, A. Annoyed users: Ads and ad-block usage in the wild. In *Proc. of IMC* (2015).
27. Ruamviboonsuk, V., Netravali, R., Uluyol, M., and Madhyastha, H. V. Vroom: Accelerating the mobile web with server-aided dependency resolution. In *Proc. of ACM SIGCOMM* (2017).
28. Sivakumar, A., Narayanan, S. P., Gopalakrishnan, V., Lee, S., Rao, S., and Sen, S. PARCEL: proxy assisted browsing in cellular networks for energy and latency reduction. In *Proc. ACM CoNEXT* (2014).
29. Sood, A. K., and Enbody, R. J. Malvertising exploiting web advertising. *Computer Fraud & Security 2011*, 4 (2011), 11 – 16.
30. Sundaresan, S., Magharei, N., Feamster, N., Teixeira, R., and Crawford, S. Web performance bottlenecks in broadband access networks. In *Proc. of SIGMETRICS* (2013).
31. The Chromium Projects. Spdy: An experimental protocol for a faster web. `https://www.chromium.org/spdy/spdy-whitepaper`, 2012.
32. The EasyList Authors. Easylist. `https://easylist.to/`, 2018.
33. Varvello, M., Blackburn, J., Naylor, D., and Papagiannaki, K. Eyeorg: A platform for crowdsourcing web quality of experience measurements. In *Proc. ACM CoNEXT* (2016).
34. Walls, R., Kilmer, E., Lageman, N., and McDaniel, P. D. Measuring the impact and perception of acceptable advertisements. In *Proc. of IMC* (2015).
35. Walton, P. User-centric performance metrics. `https://developers.google.com/web/fundamentals/performance/user-centric-performance-metrics`, 2018.
36. Wang, X. S., Balasubramanian, A., and Wetherall, D. Speeding up web page loads with shandian. In *Proc. of USENIX NSDI* (2016).
37. Wills, C. E., and Uzunoglu, D. C. What ad blockers are (and are not) doint. In *Fourth IEEE Workshop on Hot Topics in Web Systems and Technologies* (2016).
38. Wired. The new chrome and safari will reshape the web. `https://www.wired.com/2017/06/new-chrome-safari-will-reshape-web/`, 2017.
39. Work, S. How loading time affects your bottom line. `https://blog.kissmetrics.com/loading-time/`, April 28, 2011. Accessed on 05.22.2018.
40. Zarras, A., Kapravelos, A., Stringhini, G., Holz, T., Kruegel, C., and Vigna, G. The dark alleys of madison avenue: Understanding malicious advertisements. In *Proc. of IMC* (2014).