# From User Insights to Actionable Metrics: A User-Focused Evaluation of Privacy-Preserving Browser Extensions

Ritik Roongta
New York University
Brooklyn, USA
ritik.r@nyu.edu

Rachel Greenstadt
New York University
Brooklyn, USA
greenstadt@nyu.edu

## ABSTRACT

The rapid growth of web tracking via advertisements has led to an increased adoption of privacy-preserving browser extensions. These extensions are crucial for blocking trackers and enhancing the overall web browsing experience. The advertising industry is constantly changing, leading to ongoing development and improvements in both new and existing ad-blocking and anti-tracking extensions. Despite this, there is a lack of comprehensive studies exploring the set of user concerns associated with these extensions. Our research addresses this gap by identifying **five user concerns** and establishing a privacy and usability topics framework, specific to privacy-preserving extensions.

Also, many of these user concerns have not been extensively studied in the prior works. Therefore, we conducted an extensive literature review to identify shortcomings in the current benchmarking methodology. This led to the development of new techniques, including experiments to measure newly identified metrics. Our study reveals **eight new metrics** for privacy-preserving extensions that have not been previously measured. Additionally, our study enhances the measurement methodology for **two metrics**, ensuring precise results. We focus particularly on metrics that users commonly encounter on the web and report in Chrome web store reviews. Our goal is to serve as a foundational reference for future research in this field.

## CCS CONCEPTS

• **Security and privacy → Human and societal aspects of security and privacy**; **Usability in security and privacy**; *Privacy protections*;

## KEYWORDS

Web Measurement, Privacy, AdBlockers, User Concerns, Metrics, Filterlists

**ACM Reference Format:**
Ritik Roongta and Rachel Greenstadt. 2024. From User Insights to Actionable Metrics: A User-Focused Evaluation of Privacy-Preserving Browser Extensions. In *ACM Asia Conference on Computer and Communications Security (ASIA CCS '24), July 1–5, 2024, Singapore, Singapore.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3634737.3657028

## 1 INTRODUCTION

The web has become a key part of our daily lives, especially post CoVID-19. As of 2023, the average daily time spent online has surpassed six hours [16]. Reports from the World Bank and the National Institute of Health have highlighted the steep rise in e-commerce [11], health [14], and banking [9] traffic. Greater internet usage of these services leads to increased exposure of Personally Identifiable Information (PII). Previous research [45, 50, 56, 82] has shown that tracking through advertisements is not only intrusive but also deteriorates the overall web browsing experience for users. To safeguard themselves from increased web tracking and enhance their web browsing experience, users have embraced privacy-preserving extensions for their web browsers, many of which operate by blocking ads or trackers. As reported in the recent Pagefair report [8], the number of users blocking ads on the internet exceeded 820 million in 2022. Another study found that 22.2% of all users and 30% of Chrome users were using AdBlock Plus [72]. Furthermore, the FBI recommended consumers and businesses use ad-blockers in a 2022 advisory [10].

Privacy-preserving extensions typically function by implementing various features that protect against tracking, data collection, and privacy-invasive practices used by websites and advertisers. Apart from being popular, the space of privacy-preserving extensions is highly competitive. There are more than four privacy-preserving extensions with over 10 million downloads on the Chrome Web Store alone. This number rises substantially for extensions having over 100k downloads, indicating that there is no single widely accepted extension.

Using browser extensions comes with a few challenges. First, these extensions operate with the same privileges as the browser itself, creating new avenues for exploitation [58]. Second, they consume significant computational resources as they continuously check for intrusive ads and trackers. Third, extensions often need to communicate with various third parties to sync data, potentially introducing new channels for data leaks [12]. Lastly, while attempting to block ads, an extension may inadvertently break the functionality of a website, posing a usability trade-off for users who still need access to the desired content [67]. Users and their advocates face a fog of uncertainty about the impact that using privacy-preserving extensions and ad-blockers will have on their experience and how much of the differences, such as broken layouts, are intrinsic to the technologies or based on a differential treatment by sites and services.

*ASIA CCS '24, July 1–5, 2024, Singapore, Singapore*

ACM ISBN 979-8-4007-0482-6/24/07…$15.00
https://doi.org/10.1145/3634737.3657028

**Table 1: List of popular privacy-preserving extensions showcasing the diversity of the domains covered. Extensions have been categorized based on their descriptions on the web store and public perception. The popular abbreviations for each extension, developer information, number of reviews, release year, and blocking strategy for each extension are tabulated.**

| Category | Extension | Abbr. | Developer | No. of Reviews | Release year | Blocking strategy |
|---|---|---|---|---|---|---|
| Ad-Blockers & Privacy Protection | AdBlock Plus | ABP | Eyeo GmbH | 18156 | 2011 | Filterlist |
| | UBlock Origin | UbO | Raymond Hill | 9896 | 2014 | Filterlist |
| | Adguard | AdG | Adguard Software Ltd | 5610 | 2013 | Filterlist |
| | Ghostery | - | Ghostery GmbH | 3364 | 2012 | Filterlist |
| Privacy Protection | Privacy Badger | PB | Electronic Frontier Foundation | 542 | 2014 | Filterlist |
| | Decentraleyes | - | Thomas Rientjes | 107 | 2017 | CDN list |
| | Disconnect | - | Casey Oppenheim | 648 | 2011 | Filterlist |

The rapid expansion of the advertising industry has led to increasingly aggressive ad tracking and targeting techniques [24, 70]. To counter it, there has been a continuous enhancement of ad-blocking extensions. However, our preliminary study highlighted user dissatisfaction with these tools where 31% of all reviews on the Chrome web store mention at least one critical aspect of these extensions. This indicates a gap in extension developers' understanding of user needs. Additionally, many extensions fail to meet all user requirements. Therefore, it's crucial to gain a deeper insight into user concerns and identify the shortcomings in existing evaluation methods. This will enable future developers to create effective and user-centric ad-blocking extensions.

Our work addresses the following research questions:

- **RQ1:** What are the various user concerns around privacy-preserving browser extensions?
- **RQ2:** What are the unidentified metrics in the usability and privacy benchmarking methodologies of these extensions?
- **RQ3:** How to measure the novel metrics to provide a exhaustive benchmark for evaluation of the extensions?

To address these questions, we adopt a three-phased approach as shown in Figure 1. The first phase is to develop a usability and privacy framework, generated from the critical user reviews. If a user expresses discontent with any functionality of the extension in the review, we consider it as a critical aspect and classify it as a critical review. This framework contains 11 broad categories – block, ads, break, tracking, manual, filter, configuration, privacy policy, compatibility, data, and performance. We select critical reviews, using a critical score classifier, as they tend to be more informative about the problems faced by users. Our hypothesis is backed by a study done by researchers from Colorado State University [7] who found that negative information gives you more cues as compared to the expected positive information to make a decision. Each broad category further comprises a set of *related keywords* to enhance understanding and provide instructions for building new metrics. We use this framework to gauge the concerns users encounter while using these extensions. We identify and address **five major user concerns** (UCs) – *Performance, Web compatibility, Data and Privacy Policy, Extension effectiveness, and Default configurations.*

In the second phase, we perform topic modeling on the critical review dataset to pinpoint important areas for extension evaluation. We conduct a thorough literature review to determine the metrics covered by existing benchmarking methods along with their specific measurement techniques, as illustrated later in Figure 3. Our analysis reveals that, of the **14** metrics identified, **10** have not been thoroughly measured or analyzed in past studies.

In the third phase, we conduct a series of experiments to evaluate the performance of the extensions across each identified metric. We analyze each metric in-depth and offer reasonable proxies for the challenging metrics. For example, it is hard to measure reproducible breakages due to the dynamic nature of websites and the stochastic performance of automated crawls [78]. Nevertheless, our approach offers a significant step forward in understanding these metrics, laying a foundation for future, exhaustive evaluations in this domain. Our methods include techniques like web crawling, static analysis, and file comparisons. Detailed information on these methods is available in Section 5.

To conduct our study, we focus on the seven most popular privacy-preserving extensions spanning two broad categories (listed in Table 1). We select them based on their popularity as reported by AmIUnique, Wired, and MyIP [3, 13, 15][1]. Ad-Blockers & Privacy Protection (Category 1) include extensions that block ads and 3rd party trackers. Privacy Protection (Category 2) consists of extensions with the main goal of enhancing user privacy by blocking trackers. Extensions like NoScript and ScriptSafe also enhance the privacy of users but block all the JavaScript present on a page trading off usability for privacy and security. Hence, we do not use them in our analysis. Their performance can only be measured and compared when used with highly curated allowlists.

Our paper has four key contributions in the space of privacy-preserving extensions:

- We develop a fine-tuned classifier to identify and extract user reviews with critical insights.
- We generate a comprehensive framework for understanding the usability and privacy concerns of users around these browser extensions.
- We propose a comprehensive set of metrics to analyze these extensions and highlight the state of current research to identify gaps.

[1] Adblock was not included as Eyeo GmBH, Adblock Plus parent company, acquired AdBlock, Inc in April 2021

- We design new proxies to measure the performance of these extensions on the novel metrics and improve a few existing ones.

Together, these contributions provide a comprehensive taxonomy of user concerns, identify gaps in existing research, and enable effective comparison of privacy-preserving extensions.

## 2 BACKGROUND

**Browser extension** or browser plugin is a small software application that enhances the capacity or functionality of a web browser [2]. A browser extension leverages the same Application Program Interfaces (APIs) available to the website's JavaScript, in addition to its own set of APIs, thus offering additional capabilities. Privacy-preserving extensions offer advanced functionality and use sensitive permissions, making them an attractive target for adversaries. Previous studies [40, 64–67, 76, 81, 86] have extensively examined the privacy and performance of a subset of these extensions using *static analysis* and *web measurement*. Our research focuses on a superset of these studied extensions, filtering out the less popular, deprecated, or acquired ones.

**Review Analysis** is a widely used technique for understanding concerns in various domains and extracting useful features. Sentiment Analysis is often employed to aid in the process of analyzing reviews. Hu and Liu [51] studied customer reviews, using a small dataset of opinion words to calculate the sentiment of the reviews. Vu et al. [84] developed MARK, a keyword-based framework for semi-automated review analysis, employing a curated keyword dataset tailored to specific categories. A similar approach was adopted by Jindal and Liu [54]. All these approaches highlight the importance of having an initial keyword dataset to conduct such analyses which is lacking in the case of privacy-preserving extensions. Pang et al., Kanojia, and Joshi [55, 57, 71] highlighted the challenges in performing sentiment analysis on user reviews from different domains and providing plausible solutions. Nisenoff et al. [67] used user reviews and ratings to study issues with Chrome extensions.

**Topic Modeling** using Latent Dirichlet Allocation (LDA) [39] has been extensively used in unsupervised categorization of large texts, social media content, user reviews, and more. Topic classification for privacy-preserving extensions has two major challenges. First, the absence of a privacy-focused taxonomy about extension usage limits the use of supervised learning. Second, unsupervised learning generates noise due to the randomness of reviews, resulting in a wide range of irrelevant topics. To address these challenges, one can combine manual coding and analysis with LDA to enhance precision [41, 61, 69, 80]. Xue et al. [85] and Sokolova et al. [77] used LDA with qualitative analysis for topic modeling on Twitter data. Hu et al. [52] highlighted the limitations of LDA and manual analysis in identifying customer complaints through hotel reviews.

**Selenium** [32] and **Puppeteer** [31] are leading automated crawlers for web measurement and testing. Selenium is versatile, supporting multiple languages and browsers, while Puppeteer, designed for NodeJS and Chrome, excels in speed and advanced browser control features. OpenWPM [47], another framework, is built upon Firefox and Selenium and has been widely used in various studies for web and privacy measurements.

**Filterlists** serve as an extensive, independent directory of filters and host lists designed to block advertisements, trackers, malware, and various online annoyances [23]. These lists are diligently maintained by a group of dedicated researchers who regularly update them by adding new trackers and removing obsolete entries. This involves adding new trackers that emerge daily and removing those that become redundant. Some of the well-known filterlists [22] include EasyPrivacy, EasyList, Fanboy, and Peter Lowe's serverlist [30]. Importantly, these filterlists are tailored to address the multilingual nature of the web, offering support for various languages to effectively handle content from different parts of the world. Most modern ad-blockers leverage these filterlists to function, using them as a foundation to block unwanted content and enhance the user's web browsing experience.

## 3 IDENTIFYING USER CONCERNS

To identify the user concerns around privacy-preserving Chrome extensions (Table 1), we extract user reviews from Chrome web store pages of the extensions and subsequently extract usability/privacy themes from the review dataset to generate the privacy and usability topic framework (refer Table 5). Later, we identify five major user concerns from the topic framework. Only publicly available information was scraped and no personal data or user profiles were collected or stored in this process. We also got an exemption from the IRB to get a response from the developers of these extensions about our analysis of their tools. In the following subsections, we highlight the methodology used and the topics extracted.

### 3.1 Review Analysis

To gain valuable insights into users' perceptions, we analyze the reviews posted by users on the Chrome web store for the selected extensions. We scraped the reviews and support queries for the extensions in September 2022 using Selenium with a Chrome headless browser in incognito mode. By studying user perceptions, we gain insights into the usability and privacy topics that concern users when using these extensions This process involves two steps – filtering unimportant reviews from the review dataset and applying topic modeling on the refined dataset.

**Filtering Review Dataset.** We collected over 40k user reviews and support queries of the seven extensions (see Table 1) from the Chrome web store. It's essential to refine the review dataset to minimize noise in topic modeling, ensuring a focused and streamlined thematic framework. Historically, prior studies [67, 83] have relied on star ratings, often sidelining higher-rated reviews under the assumption that they depict satisfied users. However, our preliminary investigation contradicts this notion. We discovered numerous cases where users mentioned concerns about the product despite awarding it with 4 or 5 star ratings. For example, the review: *"Youtube's been updating their stuff and it stopped working at least in my region. Hope the team gets a fixed or work around soon"* highlights extension's inability to block YouTube ads despite getting 5 stars.

Adopting the previous approach might lead to ignoring the issues highlighted in these higher-rated reviews. To address this, we categorize the reviews as critical or non-critical. A review is labeled as critical if the reviewer is unsatisfied with at least one

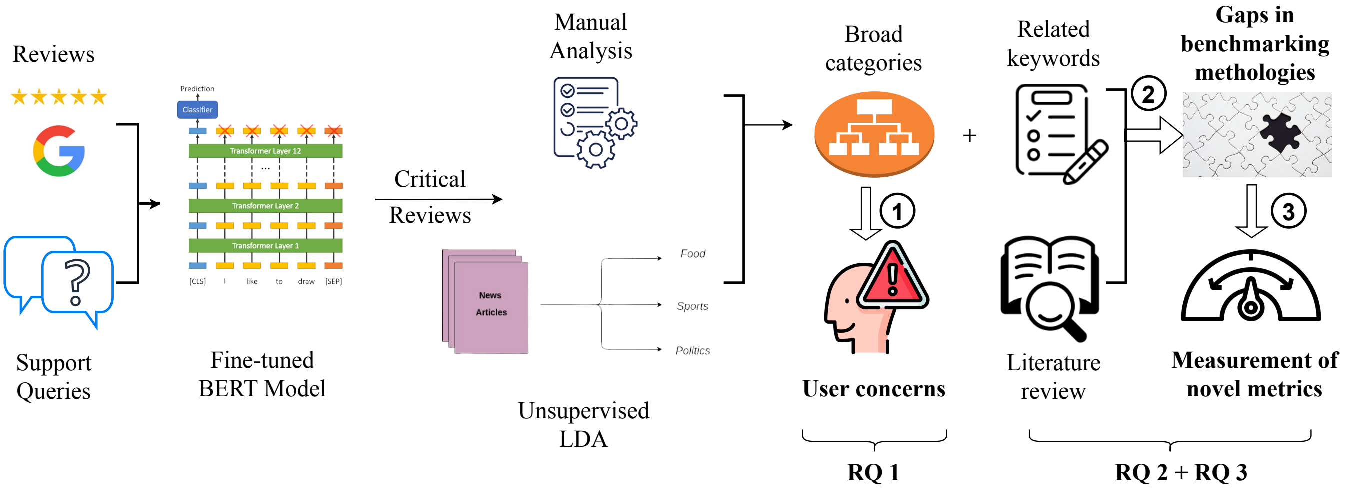

**Figure 1: The analysis unfolds in distinct stages as outlined in the flow diagram. Phase ① involves developing the topic framework using review analysis, to identify user concerns (RQ 1). In phase ②, through topic modeling and literature review, we find gaps in the benchmarking methods and introduce novel metrics for evaluation(RQ 2). Finally, phase ③ involves designing measurement experiments to evaluate the extension against the novel and existing metrics (RQ 3). The contributions are highlighted in bold.**

privacy/usability feature of the extension, even if they are generally content with its overall performance. For example, consider the review: "*Awesome until your web page recognizes it =(*". A non-critical review means that the reviewer did not express any hostile opinion about any aspect/feature of the extension. For example, review: "*It's sooo good as it blocks every ads 100% recommend this!!!*". Topic-based examples are given in Table 5 of the Appendix.

We develop a classifier to identify critical reviews by fine-tuning the BERT classifier. This technique broadly aligns with aspect-based sentiment analysis (ABSA) [79]. This classifier-driven automated tool helps to filter out non-critical reviews. To build and test this classifier, we created a training and testing dataset of 620 and 130 manually annotated reviews respectively. One author along with an undergraduate researcher independently annotated this pool of 750 reviews as either critical (assigned 0) or non-critical (assigned 1). To remove bias, the annotators reviewed the annotations, and any discrepancy (around 10%) in opinion was discussed and addressed. 45% of these reviews were annotated as critical and 55% as non-critical. The decision to use critical reviews is based on the notion that people tend to be descriptive when expressing criticism [7]. We ensure that there are enough reviews in the training dataset from every extension to cover the dynamic critiques of users across different extensions. We fine-tuned the Hugging Face sentiment analysis transformers model (**distilbert-base-uncased**)[2] by further training it on our manually annotated training dataset, thus providing us with a *criticality score*. A positive criticality score signifies that the reviewer didn't express any feature/aspect-based critique in the review. The resulting model achieved an accuracy of 86.15% on our manually annotated test dataset, surpassing the performance of the original DistilBERT model, which attained only 42.2% accuracy.

The robustness of our model is demonstrated in Figure 2. It shows the number of reviews within each criticality score bin of size 0.1. Our fine-tuned model classifies 92.33% of the reviews as critical/non-critical with a confidence level exceeding 85%. This metric increases to 95.24% for a confidence level exceeding 75%. Thus, our fine-tuned classifier differentiates between critical and non-critical reviews with high confidence. The non-critical-to-critical ratio is highest for ABP and AdG, while for UbO, it is almost equal to one. This finding is surprising, considering UbO's popularity among privacy-conscious individuals. One possible explanation is that privacy-aware users write descriptive reviews, resulting in a higher number of critical reviews for UbO. After filtering and removing out reviews with a criticality score greater than -0.7, we get a pool of 12,572 critical reviews.

## 3.2 Topic Modeling

To identify broad privacy and usability categories, we qualitatively analyze the critical reviews dataset. First, we use Latent Dirichlet Allocation (LDA) [63] along with a manual literature survey to perform unsupervised topic classification to get the initial set of codes. The literature survey consisted of an extensive exploration of various sources including privacy policies, web store descriptions of extensions, critical review dataset and support queries, and relevant peer-research literature on these extensions. These codes are further refined through manual classification to identify broad topics wherein an iterative, inductive theory-driven data coding and analysis framework is used [42]. In this manual modeling process, we streamline the noise, common in LDA, by eliminating unnecessary codes and grouping similar ones. This leads to the formation of a smaller code set, refined through iterative discussions among authors. We identify 11 broad, independent topics related

[2] https://huggingface.co/distilbert-base-uncased

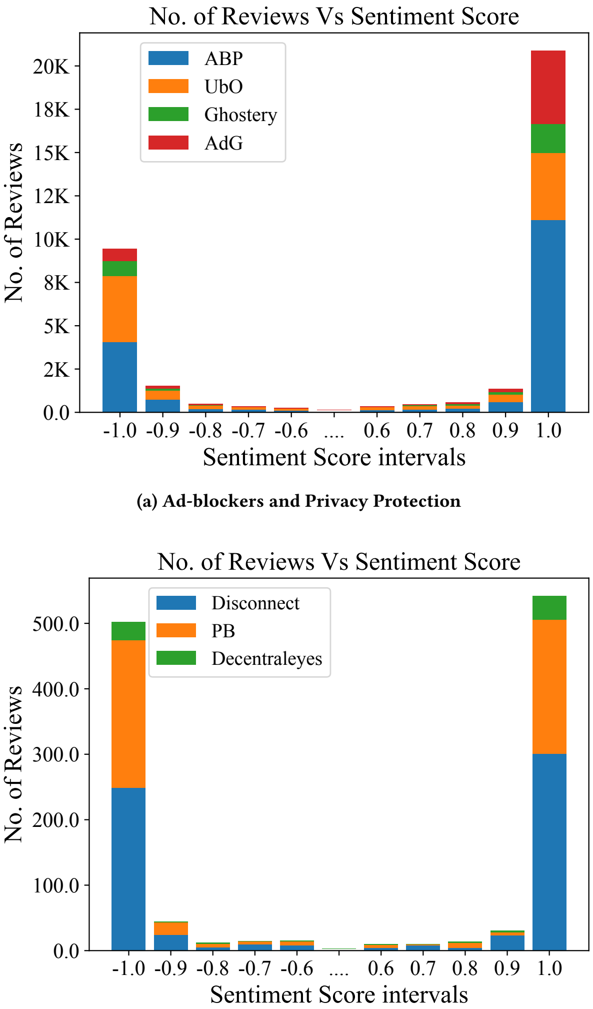

**(a) Ad-blockers and Privacy Protection**

**(b) Privacy Protection**

**Figure 2: Number of reviews within each 0.1 critical score interval. Our fine-tuned model classifies 92.33% of the reviews as either critical or non-critical with a confidence level exceeding 85% (see clusters around 1 and -1). Note: Since we use a (-1,1) annotation scale for the criticality score, there is no review with a sentiment score between -0.5 and 0.5. Hence that interval has been omitted.**

to user privacy and usability, under which other codes are categorized as subcategories or related keywords (Table 5). These related keywords are then used to categorize reviews into each broad topic through keyword extraction and matching [44], ensuring a precise and coherent classification. To establish the robustness of this framework, we randomly selected 20 reviews from each extension's critical review pool and manually examined them for potential new privacy and usability topics. As no new topics or keywords were discovered during this process, it affirmed the robustness and comprehensiveness of our initial topic framework.

**User concerns** From the developed topic framework, we pinpoint five key user concerns (UCs) by clustering together broad topics that share similar fundamental issues. The broad topics are highlighted in the brackets.

- **UC1 - Performance:** What are the scenarios where a privacy-preserving extension slows down system performance or consumes excessive memory, thus affecting computer hardware? (*performance*)
- **UC2 - Web compatibility:** When does an extension break websites or cause delays in their rendering? Is it due to anti-adblocking scripts deployed by different sites, or by blanket blocking of JavaScript by the extensions? (*compatibility, break*)
- **UC3 - Data and Privacy Policy:** How is data handled, in stationary as well as transient form, by the extensions? Can you trust the extension's permissions and policy landscape? (*data, privacy policy*)
- **UC4 - Extension effectiveness:** How effective are the extensions in blocking online tracking? Does their failure signify that they are malicious or simply incapable? (*ads, tracking, block*)
- **UC5 - Default configurations:** Which configurations and updates of these extensions have attracted critical attention from the users and made them wary of the privacy-preserving extensions? (*configuration, manual, filter*)

## 4 IDENTIFYING GAPS IN THE BENCHMARKS

Having identified the user concerns related to these extensions, our next step is to assess their effectiveness in addressing these issues. This will involve evaluating the extensions against a specific set of metrics that comprehensively represent the user concerns. Initially, we determine a set of metrics by referencing the related keywords in the topic framework, which encapsulate the various topics within each user concern. Subsequently, we review existing literature to ascertain which of these metrics have already been explored. This process helps us identify the metrics that remain unaddressed, forming the basis for our following analysis and measurement of these overlooked metrics.

In Figure 3, we highlight the different areas that concern people in different categories. For each user concern identified, we also indicate the proportion of critical reviews that are associated with that specific concern. The metrics that have been measured in the past are also highlighted with the literature references against them. The measurements for the newly identified metrics are covered later in Section 5. Each of the metrics colored in Blue indicates that there has been previous work addressing it. Orange-colored metrics indicated that these metrics are novel and haven't been measured before. Yellow-colored metrics are those metrics that have been evaluated before but could benefit from enhanced measurement methods. We find eight novel metrics for measurement – *RAM Usage, Disable Adblocker, Unable to load, Permissions, Privacy Policy and CSP, Ads, Filterlists, Acceptable Ads* while we improve measurement methods of two other metrics – *Data Usage* and *CPU Usage.*

**Performance** is mentioned in 6.01% of the critical reviews. We identified four crucial areas that users are concerned with – CPU

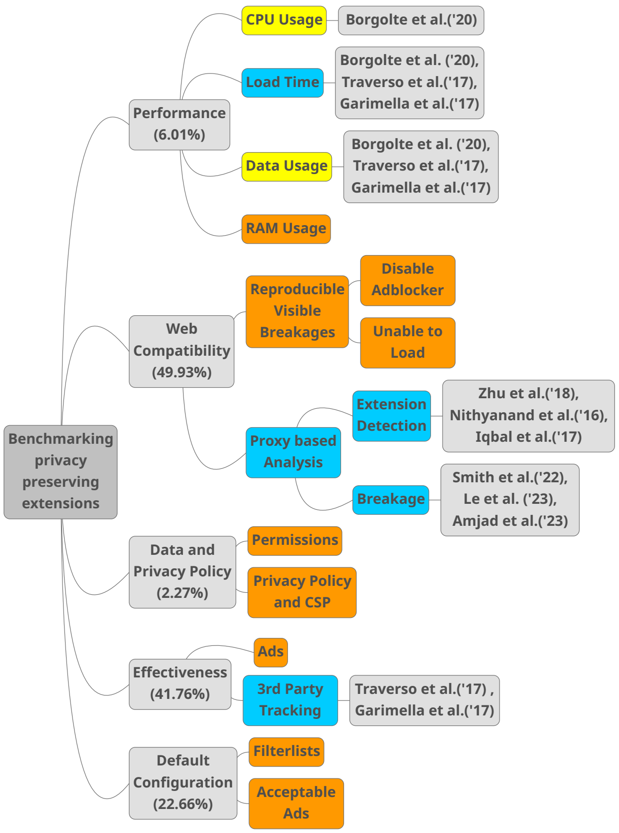
Benchmarking privacy preserving extensions
Performance (6.01%)
CPU Usage
Borgolte et al.('20)
Load Time
Borgolte et al. ('20), Traverso et al.('17), Garimella et al.('17)
Data Usage
Borgolte et al. ('20), Traverso et al.('17), Garimella et al.('17)
RAM Usage
Web Compatibility (49.93%)
Reproducible Visible Breakages
Disable Adblocker
Unable to Load
Proxy based Analysis
Extension Detection
Zhu et al.('18), Nithyanand et al.('16), Iqbal et al.('17)
Breakage
Smith et al.('22), Le et al. ('23), Amjad et al.('23)
Data and Privacy Policy (2.27%)
Permissions
Privacy Policy and CSP
Effectiveness (41.76%)
Ads
3rd Party Tracking
Traverso et al.('17) , Garimella et al.('17)
Default Configuration (22.66%)
Filterlists
Acceptable Ads

**Figure 3: A mind map illustrating the diverse User Concerns (UCs) and the corresponding metrics identified within each concern. The percentage of critical reviews discussing each concern is mentioned in brackets. For each metric, existing literature is cited where available. The metrics are visually differentiated using three colors: Blue indicates metrics that have been addressed in previous studies, Orange indicates metrics lacking prior research, and Yellow represents metrics that have been studied before but could benefit from enhanced measurement methods.**

Usage, Load Time, Data Usage, and RAM Usage through the related keywords. The most recent work in measuring performance was conducted by Borgolte et al. [40]. They used the perf tool to measure CPU usage, HAR files for Load time, and downloaded page size for Data usage on a pool of 2k live websites. Borgolte et al. [40] do not measure the RAM usage and fine-grained CPU usage. Traverso et al. [81] conducted a comparison study of a few tracker-blockers over Load time, Data usage, and the number of third parties contacted over 100 websites. Garimella et al. [49] studied Data usage, third-party tracking, and Load Time of the top 150 websites. These papers aim to analyze a selection of privacy-preserving extensions using a predefined set of metrics, focusing on a limited number of websites.

A significant gap in the prior work has been the ignorance of the upward shift of website content to replace the ad content to fill the virtual display window leading to almost similar data usage in the extension and the control case. We improve this methodology by capturing every network packet fetched for the entire scrollable page length. This accounts for the upward shift of page content after ad-blocking as well as the lazy loading [27, 28] of ads which is a common technique used for page speed enhancement. DebugBear [5] conducted a high-level study on various ad-blockers in 2021, utilizing their proprietary tools to analyze their performance on two specific websites: The Independent and the Pittsburgh Post-Gazette. They focused on metrics such as on-page CPU usage, the number of network requests, browser memory usage, and the download size of web pages. Although their analysis tools are not publicly available and their study was limited to just two websites, we observe a similar pattern between their findings on two websites and our results on 1500 tranco websites.

For **web compatibility**, we find that users most often write about cosmetic breakages in different HTML elements that can be visually observed. Web breakages are mentioned in 49.93% of all the critical reviews. Since these visual breakages are hard to measure due to the dynamic nature of the websites, we focus on two categories of breakages – 'disable adblocker' prompt detection and 'failed to load' website. Our findings are in agreement with the work done by Nisenoff et al. [67], who built a website breakage taxonomy using web store reviews and GitHub issue reports, subsequently verifying it with user experiences. According to them, unresponsiveness and extension detection (the metrics we measure) constitute 43.5% of all breakages.

Different researchers have tried to measure breakages via various proxies without explicitly reproducing those breakages. Amjad et al. [37] manually annotated 383 websites for different kinds of breakages with NoScript and UbO. They used this dataset to validate their findings about the plausible impact of blocking functional javaScripts on web breakages. Smith et al. [76] developed a classifier based on breakages caused by individual filterlist rules to predict potential breakages. They used issue reports from Easylist to train their classifier. Le et al. [59] used the change in the number of images and text on the website after installing an extension to argue about the possibility of visible breakages. A common research gap in all these studies is the lack of large-scale measurement of reproducible web breakages. For detection, various researchers have used the presence of anti-adblocking scripts on websites to argue about the potential of websites in detecting ad-blocking [53, 68, 88] but they do not measure the likelihood of a specific extension to get detected.

Our paper is the first to study **permissions, privacy policies**, and the impact of default configurations on user experience in detail for privacy-preserving extensions. Data and Privacy policy is talked about in 2.27% of critical reviews. Previous works by Carlini et al. [43] study the overall extension permission system in detail with a specific focus on security. Felt et al. [48] study the permission landscape of Android apps. Liu et al. [62] and Sanchez et al. [74] study the general browser extension landscape with the principle of least privilege (PoLP) in perspective and highlight vulnerabilities due to its violation. Although these papers do not specifically compare privacy-preserving extensions on permission abuse, they do highlight the criticality of the permissions in our evaluation set.

Around 22.66% of critical reviews express concerns about the **default configurations** in ad-blockers, a significant issue not extensively explored by previous researchers. Our analysis goes beyond comparing the filter lists utilized by different ad-blockers; we delve into the specifics of URL-based, HTML-based, and exception filter rules used by each. Additionally, we examine the rules of ABP's Acceptable Ads list and contrast them with those of other extensions. This approach provides a detailed understanding of ad-blocker configurations and their implications.

For the **effectiveness** of these extensions, users are primarily focused on ads and 3rd party trackers being blocked. It appears in 41.76% of the critical reviews. Roesner et al. [73] devised techniques for detecting the tracking abilities of third-party websites using their behavior. Siby et al. [75] developed an ML classifier for identifying tracking behavior robust to tracking and detecting adversaries. Traverso et al [81] and Garimella et al. [49] also measured the number of third parties contacted on a small set of websites by measuring the HTTP requests. We perform a high-level analysis to understand the blocking efficiency of privacy-preserving extensions rather than focusing on tracker behavior.

In the following section, we highlight the measurement methods and the subsequent results.

## 5 MEASUREMENTS

We formulate our measurement strategy to measure the new metrics identified in Section 4. To gather data for comparison, we use Selenium [32] and Puppeteer [31] based crawlers to visit websites. It is important to note that all measurements are conducted using extension versions dated May 27, 2023.

**Website testing pool.** To capture the full extension activity, it is important to evaluate the extensions on both the landing pages and inner pages of websites. Inner pages often contain more content, scripts, and ads compared to the landing page, making them valuable for our experiments. While Hispar [38], a database of internal pages, could be useful, it does not specifically focus on inner landing pages with high content. To build a suitable website testing pool, we select 1500 websites from the Tranco list [60]. This pool contains the top 1000 websites from the Tranco list and then one website for every next 500 websites, allowing us to capture a broader spectrum of websites, as lower-ranked websites may exhibit different behaviors compared to higher-ranked ones. This website pool is referred to as *basic testing pool.* We ensure that only one website (the most popular) belonging to a unique second-level domain is included, as the base policies generally do not change within the websites of the same organization. For example, we omit *www.google.co.in* since we already included *www.google.com.*

After filtering out unreachable websites we extract the three longest same-origin href links from the websites that have more than 10 href links in total. This process leaves us with our *inner page testing pool* consisting of 476 websites without any inner pages and 834 websites with three inner pages. By choosing the longest href strings, we increase the likelihood of selecting actual inner pages rather than landing pages' sections. We filter out websites with less than 10 href links on their landing page, as they tend to have only promotional or informational links and lack a lot of inner pages. Although this methodology for finding the inner pages of a website is not ideal, it provides us with a fair number of inner pages. The threshold of 10 href links is qualitatively decided by observing a sudden drop in the number of websites with more than 10 href links, followed by a consistent pattern afterward.

In the following subsections, we discuss the measurements conducted, the results obtained, and how our observations compare with the prior work. We use the Google Chrome browser for our experiments because of its popularity (market share of over 60% as of 2023). All measurements are done within docker containers, on an AMD EPYC 7542 128-core machine from a vantage point in the USA. We perform all experiments on Chrome version 113 for reproducibility. All webpages are visited three times to account for DNS caching and the average result for the metric is reported.

### 5.1 UC1: Performance

Performance is a crucial metric when it comes to evaluating browser extensions as it's discussed in over 750 critical reviews. One of the user reviews for example highlighted:

> *"Consumes way too much ram and processing power, even while browsing websites that contain zero ads."*

To measure performance, we focus on the following metrics: *CPU usage, Data usage*, and *RAM usage*. Although the first two metrics have been studied beforehand, we propose an improved approach to measure them and hence report our findings. Website *Load time* is a crucial metric; however, we do not assess it in our study as it has been effectively measured in previous research [40, 81].

For measuring *CPU usage*, the containers are configured to run on a single core, allowing us to accurately monitor CPU performance. We employ the 'mpstat' functionality instead of the perf tool utilized in the work by Borgolte et al. [40] to measure CPU usage. Instead of providing overall clock time, 'mpstat' provides us with the CPU occupancy in the user and kernel space for the duration of the process averaged over the given tick size of one second. We collect these statistics for the duration of the website load time plus an additional two seconds to allow the CPU to cool down after browser closure. Since these extensions do not spend significant time in the kernel space, we report the difference between the average value for the overall CPU usage from the control case (or no extension case) for each extension.

To measure *Data usage*, we set up a browsermob proxy [17] to intercept all request and response packets while accessing the inner page website pool via Selenium. We calculate the average of all the 'Content-Length' header fields for every website in the inner page testing pool across all inner pages. We then report the difference between the data usage measured with the extension enabled and the data usage in the control case. To account for the upward shift of content and lazy loading of ads [27, 28], we visit a web page and scroll down to the end of the page at a constant speed to capture all possible network requests from the web page.

For *RAM usage* measurement, the websites in the basic testing pool are run inside docker containers with a restricted memory limit of 4GB. We extract the RAM usage data from the 'docker stats' functionality provided by the host machine every 1.5 seconds and report the difference of average values observed during the load time for each extension from the control case.

**Table 2: Data for multiple metrics for every extension. Mean and corresponding standard deviation across all websites is shown for each extension in the performance, frames, and third-party metric. For permissions: ○ denotes not requested, ◑ denotes present but not justified i.e. the permission should not be requested or an alternative less powerful permission should be used, and ● represents that the permission is requested in a justified and reasonable way. For policy, ⊗ denotes that the policy is absent, ✓ denotes the policy does not address all the GDPR concerns as covered in PrivacyCheck, and ✔ represents either the policy complies with the GDPR concerns as covered in PrivacyCheck or claims to not collect any data.**

| | Performance | | | Break | | Storage | | Frames | 3rd Party | Permissions | | | | | Policy | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Extension | *CPU Usage (in %) (Mean (SD))* | *Data Usage (in MB) (Mean (SD))* | *RAM Usage (in MB) (Mean (SD))* | *detect* | *hang* | *localStorage (in MB)* | *IndexedDB (in MB)* | *Mean (SD)* | *Mean (SD)* | *privacy* | *all-urls* | *unlimitedStorage* | *cookies* | *notifications* | *Privacy Policy* | *CSP* |
| ABP | 13.7<br>±12.2 | 0.0<br>±2.2 | 185.7<br>±115.2 | 25 | 4 | 0.084 | 19.0 | -4.8<br>±15.2 | -40.0<br>±134.0 | ○ | ◑ | ◑ | ○ | ● | ✔ | ⊗ |
| UbO | 9.8<br>±9.4 | -1.1<br>±13.8 | 24.9<br>±101.6 | 15 | 1 | 13.0 | 0.0 | -0.2<br>±3.8 | -76.6<br>±179.9 | ● | ◑ | ● | ○ | ○ | ✔ | ✔ |
| PB | 2.3<br>±10.0 | -0.0<br>±2.4 | 53.3<br>±103.6 | 18 | 3 | 1.092 | 0.0 | -7.9<br>±18.3 | -67.0<br>±170.8 | ● | ○ | ○ | ○ | ○ | ✔ | ⊗ |
| Decentraleyes | 3.8<br>±9.6 | 0.3<br>±2.3 | 38.5<br>±80.4 | 14 | 2 | 0.124 | 0.0 | -0.1<br>±3.7 | 1.7<br>±48.5 | ● | ○ | ◑ | ○ | ○ | ✔ | ⊗ |
| Disconnect | 10.0<br>±9.6 | -0.1<br>±2.5 | 132.4<br>±107.3 | 19 | 2 | 0.150 | 0.0 | -7.8<br>±18.0 | -67.2<br>±171.1 | ○ | ○ | ○ | ○ | ○ | ✓ | ⊗ |
| Ghostery | 0.4<br>±11.5 | -0.2<br>±3.9 | 90.1<br>±102.3 | 15 | 4 | 0.712 | 6.8 | -0.4<br>±5.0 | -3.4<br>±60.7 | ○ | ○ | ○ | ● | ○ | ✔ | ✓ |
| AdG | 5.2<br>±11.7 | -0.9<br>±15.3 | 119.9<br>±104.3 | 20 | 8 | 4.412 | 0.0 | -6.8<br>±18.5 | -55.9<br>±163.9 | ◑ | ◑ | ● | ● | ○ | ✔ | ✔ |

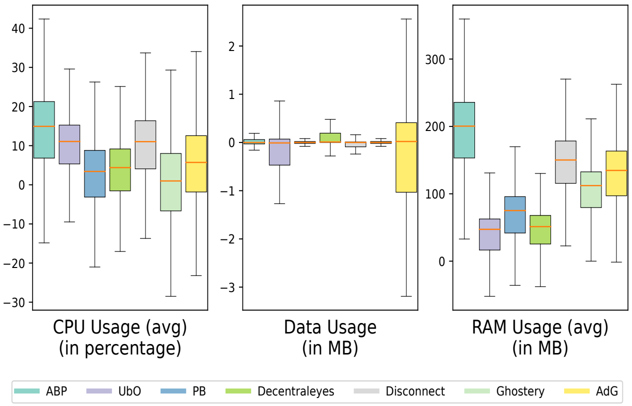

**Figure 4: Different performance metrics against every extension. The reported values are differences between the value recorded with the extension and the control case over the inner website pool. The *dotted* line represents the median of the data. Negative values signify better performance compared to the control case whereas positive median values signify poorer performance compared to the control case.**

*Results.* Figure 4 represents the performance metrics calculated for each extension using boxplots. The plotted values depict the difference between the metrics recorded when the corresponding extension is installed and the control case. Negative median values signify better performance compared to the control case whereas positive median values signify poorer performance. The outliers have not been plotted in this graph, but are depicted in Figure 6 (refer Appendix), showing the significantly high metric values for a few websites.

In terms of overall *CPU usage*, a majority of the extensions show positive median differences from the control case, indicating that they do not improve CPU performance. ABP also has a significant positive difference. PB and Ghostery perform slightly better than other extensions included in the analysis. The prior work [40] suggests that ABP performs less effectively in this metric compared to PB and Ghostery, which show similar performance levels to UbO with minimal performance enhancement. However, our granular-level findings suggest that Ghostery and PB slightly impact performance but still outperform other extensions. The variation observed in PB's performance could be linked to its comprehensive pre-training process [1], which blocks trackers from installation. This aspect of PB has seen considerable improvement in recent years.

*Data usage* exhibits a median value of almost 0 for all extensions, implying no significant difference between data transferred in the extension case compared to the control case. This might be due to additional network requests made by the ad-blockers and also websites updating their content post-blocking. Refer to Figure 6 for the outliers. Most ad-blockers have a negative average value for the difference in data usage, with UbO showing the largest negative value. AdG has the highest inter-quartile spread amongst the ad-blockers indicating that it affects data usage for a certain number of websites. ABP has a positive average value of 0.2, indicating that it transfers an additional 0.2MB of data on average.

*RAM usage* emerges as one of the most pressing concerns among users. UbO is the only ad-blocker that has the least positive median RAM usage difference, requiring an average of only 24.92 MB of additional RAM. In contrast, other ad-blockers consume significant amounts of RAM, with ABP being the worst performer, utilizing an additional average of 185.68 MB of RAM and a maximum RAM value of 315 MB.

## 5.2 UC2: Web compatibility

Web compatibility is a significant consideration when assessing the tradeoff between the security and usability of these extensions as in this review:

> *"I second the request for a quick way to 'Pause' or 'Enable/Disable' μBlock. I find that there are quite a few sites that it breaks and it is easier to just turn it off temporarily and reload the site than to try to figure out what is breaking."*

To determine a website's behavior of prompting users with a 'disable extension' dialog box, we run a Selenium crawler in virtual display mode on the inner page testing pool. This approach helps us to capture the different behaviors exhibited by websites towards extensions on their inner pages in addition to their landing pages. This discrepancy may be attributed to websites aiming to preserve ads on high-traffic and high-content pages by resisting ad filtering. During the crawling process, we visit websites, including their inner pages, both with and without the extensions. We inspect the source code of every frame for the presence of ad-blocker detector keywords such as 'adblock:detect' or 'disable ad blocker'. These keywords or their subtle variations often appear in alert boxes or iframes on the web page. Therefore, we visit each frame on the page and search for these specific keywords. We report the number of websites where such keywords are detected, indicating the websites that actively detect the presence of ad-blockers using basic JavaScript or prompt users to disable them. Additionally, we measure the number of websites that experience page hangs in the presence of extensions. A website is considered to hang if it exceeds a load time threshold of 60 seconds with the extensions enabled, while it loads within an acceptable timeframe in the control case without extensions. We selected this threshold because load times exceeding a minute are typically impractical in real-world scenarios where load times are measured in milliseconds.

*Results.* Table 2 provides an overview of the two categories of website breakages observed with the extensions. The category "detect" indicates the number of websites that either employ a basic adblock detector JavaScript or display a prompt asking users to disable their ad-blocker. The category "hang" represents the number of websites that take more than 60 seconds to load when the extensions are present. Based on our data, ABP and AdG are the most detectable extensions with 25 and 20 websites detecting them respectively. It is worth noting that ABP triggers the ad-blocker disable prompt on a significant number of websites. This finding suggests that ABP is more detectable, and websites are more willing to disable it, despite using the acceptable ads exception list.

For websites taking more than 60 seconds to load in the presence of extensions, the data reveals that there are not many instances of such behavior. AdG is associated with the highest number of websites (8) experiencing hanging or delayed loading. This number increases substantially for lower thresholds. For example, **22** websites take more than 30 seconds with ABP installed while only **2** websites take that much time with UbO. It is important to consider that this behavior is not solely dependent on the presence of extensions. Many websites have mechanisms in place to detect automated crawling, resulting in differential behavior. Manual verification of these websites was performed to ensure consistency in the data. Although the actual Load time during manual crawls varied from the automated crawls, the relative pattern remained the same i.e. websites taking longer times to load in automated crawls showed similar behavior with manual crawls.

## 5.3 UC3: Data and Privacy Policy

We identified many critical reviews about permissions, data storage, data prefetching, privacy & content security policies, and web anonymity. For example:

> *"In the latest update my chrome prompted me that the new permission "Full History Access" was added. Why?"*

While these extensions are cautious in their permission requests, even slight oversights can violate the Principle of Least Privilege (PoLP). Considering the extensive range of permissions required by these extensions, such oversights can occur frequently. To gain a deeper understanding of these issues, we conduct a qualitative analysis of the extension's permission landscape.

**Permissions.** In our evaluation of the extensions, we focus on five specific permissions: *privacy*, *unlimitedStorage*, *all-urls*, *notifications*, and *cookies*. Our goal is to determine if the permissions requested by the extensions are necessary for the tasks they perform. Regarding data storage capabilities, we explore two commonly used techniques: *localStorage* and *IndexedDB*. LocalStorage and IndexedDB are commonly used by extensions to store various data objects required for their functionality. Refer to their documentation [20, 26] for in-depth information about their functionality.

As per Chrome documentation [20], the *all-urls* permission matches any URL that starts with a permitted scheme (http:, https:, file:, ws:, wss: or ftp:). The *privacy* permission enables network prefetching, webRTC blocking, etc., and is considered sensitive. It triggers a warning message that it "can change your privacy-related settings," which has drawn attention from users and raised concerns about its usage. The *cookies* permission gives extensions access to the chrome.cookies API. The *notifications* permission enables extensions to create and show notifications.

**Privacy policy and CSP.** Some of the extensions in our study lack clearly defined privacy policies, even though they are available in the European region, which falls under the jurisdiction of the General Data Protection Regulation (GDPR) [33]. Given that these extensions process various aspects of PII associated with users, they must have well-defined privacy policies, at least in principle. To assess and compare the privacy policies of these extensions, we use the PrivacyCheck extension developed by Zaeem et al. [87]. It enables us to summarize and evaluate the privacy policies of the extensions under consideration. Additionally, we also examine the content security policies (CSP) implemented by these extensions via source code analysis.

*Results.* **Permissions.** The results of the source code examination of the extensions are shown in Table 2. *unlimitedStorage* permission is requested by four extensions but is only used by UbO and AdG. Table 2 provides information on the local storage space utilized by each extension–only UbO exceeds the 5 MB limit. Some extensions argue that this permission is necessary because users can add large filterlists that may surpass the limit. However, extensions like Ghostery and PB allow filterlists without requiring this permission, demonstrating that alternatives exist. Moreover, the inability to optionally use this permission becomes an easy justification for the extension developers to request it, even though it serves a tiny population of users who might import large filterlists.

*all-urls* permission is requested by ABP, UbO, and AdG. However, Chrome and Mozilla disabled FTP support in their respective browsers in 2021 [4]. Additionally, the limited number of "file:" based URLs are usually static so there appears to be minimal scope for these extensions to block content on those web pages. Therefore, it is recommended to use strict URL matching for "http:" and "ws:" protocols only.

During our review study, many users expressed dissatisfaction with receiving annoying and intrusive notifications. Some notifications are a result of the extensions' inability to block them, while others are self-generated by the extensions. ABP uses this permission to redirect users to their donation and review pages, respectively. Considering the critical user sentiment towards these notifications, extensions might consider providing this information in their descriptions instead. The *cookies* permission is requested by Ghostery and AdG. Ghostery requests this permission to block cookies and check for logged-in users. AdG uses it to remove cookies that match its filterlist rules. Overall, the usage of this permission by these extensions seems to be for benign purposes, aligning with their functionality and privacy objectives.

Among the four extensions that request the *privacy* permission (UbO, AdG, PB, and Decentraleyes), AdG is the only extension that lists it as an optional permission. The primary reason for requesting this permission is to disable network prefetching[3], hyperlink auditing[4], and block WebRTC requests[5].

Decentraleyes, PB, and UbO use this permission to disable prefetching. PB and UbO also use it to disable hyperlink auditing. Additionally, PB employs this permission to block alternate error pages, where setting the policy to True allows Google Chrome to use built-in alternate error pages (such as "page not found"). AdG uses this permission to block WebRTC functionality. However, UbO recently removed the WebRTC blocking option, as both Chrome and Firefox no longer leak private IP addresses [6]. Requesting this permission solely for WebRTC blocking, as done by AdG, may be unnecessary.

**Privacy Policy and CSP.** Some extensions lack a privacy policy or claim no data collection, operating solely on the client side. We used PrivacyCheck for 10 GDPR checks on the remaining extensions' policies. None inform users of data breaches. ABP and AdG restrict PII use from minors under 16. Most extensions, except Disconnect, comply with other GDPR sections covered by PrivacyCheck. Regarding the handling of PII categories, ABP complies with COPPA, allowing users to opt out in case their privacy policy changes. Disconnect lacks user data control, such as editing collected information. While all extensions collect PII like email addresses, none gather sensitive data like SSNs or credit card details. All extensions may disclose information to government authorities as legally required, such as subpoenas.

In terms of Content Security Policy (CSP) implementation, UbO, Ghostery, and AdG use CSP directives in their manifest files. These extensions configure the 'script-src' and 'object-src' fields to 'self', allowing local plugin content and script origins. Ghostery implements 'wasm-eval', which may be considered unsafe and is not included in MDN's CSP directives [19].

## 5.4 UC4: Extension Effectiveness

This user concern raises critical questions about the improvements offered by the extensions in web browsing. It encompasses two main categories: *ads* and *tracking*.

> *"double click advertising is tracking me even blocked what can i do?? Help block them! I also get a lot of Google analytics and google platforms should i trust them or what??"*

We address each category individually, assuming that the websites are rendered without breakage and that increased filtering indicates better performance. These measurements are high-level proxies to provide us insights into how many ads and third parties might have been blocked by these extensions without measuring advanced adversaries that employ extension evasion strategies. For calculating both Ads and third parties, we visit each website in the inner page testing pool, three times, and report the average difference between the extension and the control case. This approach accounts for the dynamic nature of the websites.

**Ads.** To evaluate the effectiveness of ad-blocking, we measure the reduction in the number of frames on web pages. Advertisements are often displayed within HTML frames, which display content independent of its container. Many HTML tags that contain ad scripts are rendered as frames and iframes. Effective ad-blockers remove the frame objects from the page layout as part of the filtering process. We use Puppeteer to hook into the web page and calculate the number of frames using the page.metrics() function.

**Third parties.** A 3rd-party website domain refers to the domain or web address of a website that is operated and owned by a separate entity or organization distinct from the owner/operator of the primary website. Most trackers can be classified as a 3rd party domain. OpenWPM [47] provides us with a mechanism to calculate the number of 3rd parties contacted during a website crawl.

*Results.* Table 2 shows the average and standard deviation values obtained from our data collection for each crawl. The mean value represents the average reduction in the recorded metric. A higher

[3] developer.mozilla.org/en-US/docs/Web/Performance/dns-prefetch
[4] www.thewindowsclub.com/ping-hyperlink-auditing-in-chrome-firefox
[5] developer.mozilla.org/en-US/docs/Web/API/WebRTC_API

negative mean value suggests that the extension is more effective at blocking ads and third-party content.

**Ads.** UbO and Ghostery perform the least effectively, blocking only 0.2 and 0.4 frames on average. PB performs well by blocking an average of 7.9 frames.

**Third parties.** UbO performs best, blocking an average of 77 third parties, while Ghostery has the lowest effectiveness, blocking only 3.4 third parties. Decentraleyes actually *increases* the number of third parties contacted on average.

## 5.5 UC5: Default Configurations

Default configurations are significant user concerns that encompass categories such as manual settings, filters, and configurations. For example:

> *"Great app, but it lacks default blocking options. There are services I want to block everywhere and there are services I don't want to block at all (e.g. facebook)."*

Many users found the necessity to manually configure allowlists or adjust extension settings to be bothersome. This aligns with the common understanding that users typically prefer to use tools with their default configurations [21, 29]. Quantifying this through automated testing is challenging. We examine the default filter lists used by ad-blockers and measure the different URLs and HTML elements blocked by these lists along with the exception rules.

We identify the filterlists imported by the extensions and calculate the number of third-party domains allowed by the Acceptable Ads exception list of ABP that are blocked by other extensions. The presence of Acceptable Ads is a significant concern for users, as they may encounter ads on websites despite having ad-blocking extensions installed. This discrepancy in ad-blocking effectiveness can lead to user dissatisfaction, as is evident in one of the ABP's user reviews (e.g., **Review:** *They used to block adds but now their "acceptable adds" are just as bad as what they used to block.*). We focus specifically on third-party domains allowed by the Acceptable Ads list, excluding any HTML tags or cookies.

**Table 3: Number of URL-based, HTML-based and exception filter rules used by the extensions in their default state. PB and Disconnect neither have filter rules for HTML-based elements nor Exception rules.**

| Extension | URL-based filter rules | HTML-based filter rules | Exception rules |
|---|---|---|---|
| ABP | 75992 | 293403 | 13117 |
| UbO | 43202 | 22808 | 3088 |
| PB | 2171 | - | - |
| Disconnect | 2506 | - | - |
| Ghostery | 75317 | 38249 | 5328 |
| AdG | 263912 | 117607 | 13364 |

*Results.* **Filterlists.** ABP, UbO, and Ghostery use popular filterlists such as Easylist and Easyprivacy. Ghostery and UbO also import Peter Lowe's filterlist and Fanboy, and UbO has its additional lists as well. AdG uses its own set of desktop and mobile filterlists. PB begins with a seed file to calibrate its third-party tracking algorithm. Disconnect's filterlist was last updated in 2019 before the transition to a paid service. Decentraleyes maintains a local copy of popular JS libraries to prevent tracking by popular CDN servers. Table 3 shows the number of URL-based filter tiles, HTML elements-based filter rules, and exception rules for each extension in their default mode. UBO and Ghostery have a similar number of filter rules while ABP has significantly fewer filter rules among the filterlist-based extensions. Also, ABP has a significantly high number of exception rules due to the Acceptable Ads list. AdG has the highest number of filters as well as exception rules because it uses its own big set of filterlists for blocking.

**Acceptable Ads.** Many users express concerns about the use of the Acceptable Ads campaign and the corresponding allowlist used by ABP. While other extensions also allow certain URLs for compatibility purposes, the Acceptable Ads allowlist aims to support advertisers who adhere to the Acceptable Ads Standard policies [18]. This list contains a total of 10,934 exception URLs that are allowed. To assess the treatment of these URLs by other extensions, we focus on AdG, UbO, and Ghostery since they use different filterlists to block ads. From the Acceptable Ads allowlist, we find that UbO, AdG, Ghostery block 277, 442, 293 URLs and allow 143, 258, 150 URLs respectively. The remaining URLs from the Acceptable Ads allowlist are not mentioned in the blocked or allowed sections of the respective filterlists.

# 6 SUMMARY

In our work, we initially pinpoint user concerns regarding privacy-preserving extensions by analyzing reviews from the Chrome Web Store. This analysis leads to the creation of a framework focusing on usability and privacy topics, revealing five key user concerns. From this framework, we identify 14 critical metrics for evaluating these extensions. We discover that existing literature does not address 10 out of these 14 metrics effectively. To bridge this research gap, we design measurement methods and apply them to evaluate the performance of various extensions on these metrics.

In Table 4, we summarize our findings for each metric analyzed. We focus on major metrics evaluated in this paper and mention the ideal and subpar extensions. As observed here, it is challenging for the existing pool of extensions to address all user concerns but can help them to evaluate the extensions based on their priorities. Table 4 underscores several novel metrics such as Data Usage, RAM usage, permissions, reproducible detection, unresponsiveness, etc. which are essential for assessing privacy-preserving extensions. This evaluation, encompassing a broad range of user concerns, enables individuals to make informed decisions about which extension or combination of extensions to install, based on their specific priorities and areas of interest.

# 7 DISCUSSION

**Recommendations.** We reached out to the developers of all the extensions regarding the instances of permission abuse. We got an unsatisfactory response from the developers of a few extensions citing legacy support as the reason for the continuation of using such permissions. We have the following recommendations for the developers:

**Table 4: Summary of extensions' performance on different metrics. *Ideal* represents the best-performing set of extensions, and *Subpar* represents poorly performing extensions. Abbreviations: D'nect-Disconnect, Gh'ry-Ghostery.**

| **Metrics** | *Ideal* | *Subpar* |
|---|---|---|
| CPU usage | Gh'ry | ABP |
| Data Usage | UbO | ABP |
| RAM usage | UbO | ABP |
| ad-blocker detection prompt | UbO<br>Gh'ry | ABP |
| Permissions | Gh'ry | AdG |
| Privacy Policy | UbO<br>ABP | Gh'ry |
| Ads | PB<br>D'nect | UbO |
| 3rd-party | UbO<br>PB<br>D'nect | ABP |
| Default filterlists | AdG | ABP |

- Maintain a version database to support legacy browsers while shipping the most secure version in the web store.
- Develop a sound issue reporting forum for users to report issues with the extension usage.
- Work with browser vendors to design permission APIs adhering to security and privacy principles.

In our research, we identify certain limitations and complex issues, which merit further attention from fellow researchers in the field of security and privacy. We highlight them here.

**Reviews.** We face two significant challenges with the review dataset. First, we need a method to differentiate fake reviews from genuine ones [34] and prevent the impact of review mills [35]. This issue is challenging to address as there are no accurate verification methods other than Google's proprietary machine learning techniques, while other tools like Fakespot [36] are limited to Amazon, Walmart, and eBay reviews. Second, the reviews are spread over a period of 10 years, which introduces concerns that may no longer be relevant. Also, the uneven and sporadic distribution of reviews across different periods makes it hard to track shifts in topics over time. To address this, we conduct a measurement study to determine if the concerns identified are still relevant and have a noticeable impact. Also, by considering only critical reviews, we ensure that short reviews (less than 5 words), that do not mention any aspect of the extensions and simply are used to increase overall star ratings, are filtered out. For more details on the temporal shift of critical sentiments, refer to Appendix A.

**Breakages.** Measuring website breakages deterministically is a hard problem. They depend on several factors operating in tandem, and isolating them is difficult. While Nisenoff et al. [67] offers a foundational taxonomy of breakages, there remains a gap in mapping them to their respective causes. Existing GitHub issue reports of website breakages are difficult to rely upon as popular sites are dynamic and get patched quickly. Our work tries to address a few of these challenges by building a taxonomy of types of breakages around privacy-preserving extensions and testing a few of them like detection and unresponsiveness which can be measured deterministically. Due to the dynamic nature of the websites, the breakages are hard to measure as they get patched quickly and might not appear reliably across different testing environments. Our future work focuses on detecting visible and non-visible breakages on websites using AI-assisted crawling and manual analysis.

**Fingerprint.** Fingerprinting of users in the presence of extensions is a complex topic. Although the change in the user fingerprint due to a single attribute is often studied, analyzing the effect on the entire fingerprinting landscape of a user is tough since it depends on multiple factors like the strength of the adversary, the interdependence of different attributes, etc. EFF's tool [46] considers different attributes but does not incorporate their interdependence. Future work should aim at studying the collective impact of each web attribute on the user's entire web footprint.

**Benchmarking framework.** Each user concern in our analysis is distinct and has the potential to be studied as an independent research topic. The benchmarking process does not have a standardized set of measurement methods for every concern. This is evident in our work where the depth of evaluation for each user concern varies. For example, we conduct a detailed evaluation of the 'performance' user concern due to the precise mention of various performance-related issues in user reviews. Conversely, the 'break' concern is evaluated superficially due to the vagueness of user reviews and the complexity of evaluation methodologies.

## 8 CONCLUSION

Our study emphasizes the importance of a detailed assessment of user concerns to enhance the efficacy of privacy-preserving browser extensions. We identify 14 key metrics that impact user experience with these extensions. Our extensive literature review reveals that 10 of these metrics lack proper measurement methods or could benefit from enhanced measurement techniques. We've developed new methods and improved existing ones to evaluate the performance of extensions across these 10 metrics. Our findings are summarized in Table 4, where we distinguish between ideal and subpar extensions for each metric. Additionally, we point out intricate aspects of extension benchmarking, such as website breakages, that need further exploration. This research lays the groundwork for future studies in enhancing ad-blocking and privacy protection extensions.

The centralized code repository for our project can be found here[6]. This repository includes the code for the measurements. We also release the fine-tuned BERT model[7].

## ACKNOWLEDGMENTS

We thank Mitchell Zhou for his assistance in manually annotating critical reviews. We would also like to thank Ben Stock and other anonymous reviewers for their feedback and insight.

[6] https://github.com/Racro/measurements_user-concerns

[7] https://huggingface.co/racro/sentiment-analysis-browser-extension

## REFERENCES

[1] 2018. Giving Privacy Badger a Jump Start. Online. https://www.eff.org/deeplinks/2018/08/giving-privacy-badger-jump-start

[2] 2018. What is a Browser Extension? Online. https://www.techtarget.com/whatis/definition/browser-extension

[3] 2020. These Chrome extensions protect you against creepy web tracking. Online. https://www.wired.co.uk/article/chrome-extensions-privacy-ad-tracking-blocker

[4] 2021. Google Chrome 88 removes Flash and FTP support. Online. https://hexus.net/tech/news/software/147315-google-chrome-88-removes-flash-ftp-support/

[5] 2021. How Do Chrome Extensions Impact Browser Performance? Online. https://www.debugbear.com/blog/chrome-extension-performance-2021#increasing-website-cpu-usage

[6] 2021. Remove WebRTC leak prevention. Online. https://github.com/uBlockOrigin/uBlock-issues/issues/1723

[7] 2021. Why negative reviews could have more of an impact on some of the most important customers. Online. https://biz.source.colostate.edu/negative-online-reviews-impact-study/

[8] 2022. 2022 PageFair Adblock Report. Online. https://blockthrough.com/blog/2022-pagefair-adblock-report/

[9] 2022. COVID-19 Drives Global Surge in use of Digital Payments. Online. https://www.worldbank.org/en/news/press-release/2022/06/29/covid-19-drives-global-surge-in-use-of-digital-payments

[10] 2022. Cyber Criminals Impersonating Brands Using Search Engine Advertisement Services to Defraud Users. Online. https://www.ic3.gov/Media/Y2022/PSA221221?=8324278624

[11] 2022. E-Commerce Sales Surged During the Pandemic. Online. https://www.census.gov/library/stories/2022/04/ecommerce-sales-surged-during-pandemic.html

[12] 2022. Google Chrome, Firefox ad blocker extensions leaked million users data. Online. https://tech.hindustantimes.com/tech/news/google-chrome-firefox-ad-blocker-extensions-leaked-million-users-data-story.html

[13] 2022. Privacy Tools - Am I Unique? Online. https://amiunique.org/privacy-tools/

[14] 2022. Virtual care use during the COVID-19 pandemic. Online. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9037937/

[15] 2023. 5 Browser Extensions to Protect Your Privacy. Online. https://whatismyipaddress.com/privacy-browser-extensions

[16] 2023. Average daily time spent using the internet by online users worldwide from 3rd quarter 2015 to 2nd quarter 2023. Online. https://www.statista.com/statistics/daily-time-spent-online-global/

[17] 2023. BrowserMob Proxy. Online. https://github.com/lightbody/browsermob-proxy

[18] 2023. Building bridges. Online. https://acceptableads.com/

[19] 2023. Content-Security-Policy. Online. https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Content-Security-Policy

[20] 2023. Declare permissions. Online. https://developer.chrome.com/docs/extensions/mv3/declare_permissions/

[21] 2023. Default settings. Online. https://www.marketingsociety.com/the-gym/default-settings-most-powerful-tool-behavioural-scientist’s-toolbox

[22] 2023. EasyList. Online. https://easylist.to/

[23] 2023. Filter list. Online. https://brave.com/glossary/filter-list/

[24] 2023. FTC Report Shows Rise in Sophisticated Dark Patterns Designed to Trick and Trap Consumers. Online. https://www.ftc.gov/news-events/news/press-releases/2022/09/ftc-report-shows-rise-sophisticated-dark-patterns-designed-trick-trap-consumers

[25] 2023. IAB Categories. Online. https://docs.webshrinker.com/v3/iab-website-categories.html#iab-categories

[26] 2023. IndexedDB API. Online. https://developer.mozilla.org/en-US/docs/Web/API/IndexedDB_API

[27] 2023. Lazy loading. Online. https://developers.google.com/publisher-tag/samples/lazy-loading

[28] 2023. Lazy Loading Ads: What, When, Where, Why & How. Online. https://www.adpushup.com/blog/lazy-loading-ads/

[29] 2023. New frontiers: default settings. Online. https://www.research-live.com/article/opinion/new-frontiers-default-settings/id/5052872

[30] 2023. Peter Lowe’s block list. Online. https://pgl.yoyo.org/adservers/serverlist.php

[31] 2023. Puppeteer. Online. https://pptr.dev/

[32] 2023. Selenium automates browsers. That’s it! Online. https://www.selenium.dev/

[33] 2023. What is GDPR, the EU’s new data protection law? Online. https://gdpr.eu/what-is-gdpr/

[34] 2024. How to Tell if Reviews are Fake: Spot Fake from Real Reviews. Online. https://reputation.com/resources/articles/spot-fake-reviews-how-to/

[35] 2024. Review mills identified as a new form of peer-review fraud. Online. https://www.chemistryworld.com/news/review-mills-identified-as-a-new-form-of-peer-review-fraud/4018888.article

[36] 2024. Use AI to detect fake reviews and scams. Online. https://www.fakespot.com/

[37] A H Amjad, Z. Shafiq, and M A Gulzar. 2023. Blocking JavaScript Without Breaking the Web: An Empirical Investigation. In *PoPETs*.

[38] W Aqeel, B Chandrasekaran, A Feldmann, and B M Maggs. 2020. On Landing and Internal Web Pages. In *ACM Internet Measurement Conference (IMC ’20)*.

[39] David M Blei, A Y. Ng, and M I Jordan. 2003. Latent Dirichlet Allocation. In *Journal of Machine Learning Research 3.*

[40] K Borgolte and N Feamster. 2020. Understanding the Performance Costs and Benefits of Privacy-focused Browser Extensions. In *The Web Conference.*

[41] D Bowie-DaBreo, C Sas, H. Iles-Smith, and S Sünram-Lea. 2022. User Perspectives and Ethical Experiences of Apps for Depression: A Qualitative Analysis of User Reviews. In *Conference on Human Factors in Computing Systems.*

[42] V Braun and V Clarke. 2011. THEMATIC ANALYSIS. In *APA Handbook of Research Methods in Psychology.*

[43] N Carlini, A P. Felt, and D Wagner. 2012. An Evaluation of the Google Chrome Extension Security Architecture. In *USENIX.*

[44] Gopan E, Rajesh S, Vishnu GR, Akhil Raj R, and Thushara MG. 2020. Comparative Study on Different Approaches in Keyword Extraction. In *ICCMC.*

[45] Zeng E, Kohno T, and Roesner F. 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.*

[46] P Eckersley. 2010. How Unique Is Your Web Browser?. In *PETS.*

[47] S Englehardt and A Narayanan. 2016. Online Tracking: A 1-million-site Measurement and Analysis. In *ACM CCS.*

[48] A P Felt, E Chin, S Hanna, D Song, and Wagner D. 2012. Android permissions demystified. In *ACM conference on Computer and communications security.*

[49] K Garimella, O Kostakis, and M Mathioudakis. 2017. Ad-blocking: A Study on Performance, Privacy and Counter-measures. In *WebSci.*

[50] Yung H and Oliver MB. 2004. Exploring the Effects of Online Advertising on Readers’ Perceptions of Online News. In *Journalism & Mass Communication Quarterly.*

[51] M. Hu and B Liu. 2004. Mining and Summarizing Customer Reviews. In *Knowledge Discovery and Data Mining.*

[52] N Hu, T Zhang, B. Gaob, and I Bose. 2019. What do hotel customers complain about? Text analysis using structural topic model. In *Tourism Management Journal.*

[53] Umar Iqbal, Zubair Shafiq, and Zhiyun Qian. 2017. The Ad Wars: Retrospective Measurement and Analysis of Anti-Adblock Filter Lists. In *Proceedings of the 2017 Internet Measurement Conference.*

[54] N. Jindal and B Liu. 2008. Opinion Spam and Analysis. In *International Conference on Web Search and Data Mining.*

[55] A Joshi. 2019. Sentiment Analysis and Opinion Mining from Noisy Social Media Content. In *Master’s Thesis in IIIT, Hyderabad.*

[56] Lim JS, Chock TM, and Golan GJ. 2020. Consumer perceptions of online advertising of weight loss products: the role of social norms and perceived deception. In *Journal of Marketing Communication.*

[57] D. Kanojia and A Joshi. 2023. Applications and Challenges of SA in Real-life Scenarios. In *Computational Intelligence Applications for Text and Sentiment Data Analysis.*

[58] Y M Kim and B lee. 2023. Extending a Hand to Attackers: Browser Privilege Escalation Attacks via Extensions. In *USENIX.*

[59] H Le, S Elmalaki, and A Markopoulou. 2023. AutoFR: Automated Filter Rule Generation for Adblocking. In *USENIX.*

[60] V Le Pochat, T Van Goethem, S Tajalizadehkhoob, M. Korczy ´nski, and W Joosen. 2019. TRANCO: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *NDSS.*

[61] S.E Levy, W. Duan, and S Boo. 2013. An analysis of one-star online reviews and responses in the Washington, DC, lodging market. In *Cornell Hospitality Quarterly.*

[62] L Liu, X Zhang, G Yan, and S Chen. 2012. Chrome Extensions: Threat Analysis and Countermeasures. In *NDSS security.*

[63] Blei D M., Ng A Y., and Jordan M I. 2003. Latent Dirichlet Allocation. In *Journal of Machine Learning Research 3.*

[64] A Mathur, J Vitak, A Narayanan, and M Chetty. 2018. Characterizing the Use of Browser-Based Blocking Extensions To Prevent Online Tracking. In *USENIX Symposium on Usable Privacy and Security (SOUPS).*

[65] J Mazel, R Garnier, and K Fukuda. 2019. A comparison of web privacy protection techniques. In *Computer Communications.*

[66] G Merzdovnik, M Huber, D Buhov, N Nikiforakis, S Neuner, M Schmiedecker, and E Weippl. 2017. Block Me If You Can: A Large-Scale Study of Tracker-Blocking Tools. In *IEEE European Symposium on Security and Privacy.*

[67] A Nisenoff, A Borem, M Pickering, G Nakanishi, M Thumpasery, and B Ur. 2023. Defining "Broken": User Experiences and Remediation Tactics When Ad-Blocking or Tracking-Protection Tools Break a Website’s User Experience. In *Usenix Security Symposium.*

[68] R Nithyanand, S Khattak, M Javed, Narseo Vallina-Rodriguez, Marjan Falahrastegar, JE Powles, Emiliano De Cristofaro, Hamed Haddadi, and Steven Murdoch.

2016. Adblocking and Counter Blocking: A Slice of the Arms Race. In *6th USENIX Workshop on Free and Open Communications on the Internet (FOCI '16)*.

[69] J Nitkowski. 2022. Qualitative analysis of user reviews from Nurx and Planned Parenthood Direct: what user experiences reveal about telecontraception apps. In *Sexual Health, Collingwood Vol. 19*.

[70] Agarwal P, Joglekar S, Papadopoulos P, Sastry N, and Kourtellis N. 2020. Stop tracking me Bro! Differential Tracking of User Demographics on Hyper-Partisan Websites. In *WWW*.

[71] B. Pang and L Lee. 2008. Opinion Mining and Sentiment Analysis. In *Foundations and Trends in Information Retrieval*.

[72] Enric Pujol, Oliver Hohlfeld, and Anja Feldmann. 2015. Annoyed Users: Ads and Ad-Block Usage in the Wild. In *Proceedings of the 2015 Internet Measurement Conference (IMC '15)*. https://doi.org/10.1145/2815675.2815705

[73] F Roesner, T. Kohno, and D Wetherall. 2012. Detecting and Defending Against Third-Party Tracking on the Web. In *USENIX NSDI*.

[74] P P Sanchez, L Ortiz-Martin, G Schneider, and A Sabelfeld. 2022. Are chrome extensions compliant with the spirit of least privilege?. In *International Journal of Information Security*.

[75] S Siby, U Iqbal, S Englehardt, Z Shafiq, and C Troncoso. 2023. WebGraph: Capturing Advertising and Tracking Information Flows for Robust Blocking. In *USENIX*.

[76] M Smith, P Snyder, M Haller, B Livshits, D Stefan, and H Haddadi. 2022. Blocked or Broken? Automatically Detecting When Privacy Interventions Break Websites;. In *PoPETs*.

[77] M Sokolova, K Huang, S Matwin, J Ramisch, V Sazonova, R Black, C Orwa, S Ochieng, and N Sambuli. 2016. Topic Modelling and Event Identification from Twitter Textual Data. In *Social and Information Networks (cs.SI)*.

[78] Ahmad SS, Dar MD, Zaffar MF, Vallina-Rodriguez N, and Nithyanand R. 2020. Apophanies or Epiphanies? How Crawlers Impact Our Understanding of the Web. In *WWW*.

[79] C Sun, L Huang, and X Qiu. 2019. Utilizing BERT for Aspect-Based Sentiment Analysis via Constructing Auxiliary Sentence. In *NAACL*.

[80] M L Tan, P Raj, K Stock, Emma E H. Doyle, and L Graham. 2020. Modified Usability Framework for Disaster Apps: A Qualitative Thematic Analysis of User Reviews. In *International Journal of Disaster Risk*.

[81] S Traverso, M Trevisan, L Giannantoni, M Mellia, and H Metwalley. 2017. Benchmark and comparison of tracker-blockers: Should you trust them?. In *Network Traffic Measurement and Analysis Conference (TMA)*.

[82] Blase Ur, Leon PG, Cranor LF, Shay R, and Wang Y. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *SOUPS*.

[83] S Vetrivel, V van Harten, C H Gañán, M van Eeten, and S Parkin. 2023. Examining Consumer Reviews to Understand Security and Privacy Issues in the Market of Smart Home Devices. In *USENIX*.

[84] P M Vu, T T Nguyen, H V. Pham, and T T Nguyen. 2015. Mining User Opinions in Mobile App Reviews: A Keyword-based Approach. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*.

[85] J Xue, J Chen, C Chen, C Zheng, S. Li, and T Zhu. 2020. Public discourse and sentiment during the COVID 19 pandemic: Using Latent Dirichlet Allocation for topic modeling on Twitter. In *PLoS ONE Journal*.

[86] Z Yu, S Macbeth, K Modi, and J Pujol. 2016. Tracking the trackers. In *International World Wide Web Conference*.

[87] Razieh N. Zaeem, Rachel L. German, and K S. Barber. 2018. PrivacyCheck: Automatic Summarization of Privacy Policies Using Data Mining. In *ACM Transactions on Internet Technology, Volume 18, Issue 4*.

[88] Shitong Zhu, Xunchao Hu, Zhiyun Qian, Zubair Shafiq, and Heng Yin. 2018. Measuring and Disrupting Anti-Adblockers Using Differential Execution Analysis. In *Network and Distributed System Symposium (NDSS)*.

# A CHANGE IN SENTIMENTS OVER TIME

Figure 5 illustrates the changes in the criticality score over time, with the number of critical reviews plotted on the secondary vertical axis. The sentiment analysis, covering 60-day intervals, shows no distinct trends or patterns. Due to the low number of reviews for other extensions, we limited our analysis to just five extensions. However, longer-term observations reveal some patterns. For instance, Ghostery experienced a notable decline in sentiment from 2016 to 2020, coinciding with allegations of user data misuse. Additionally, ABP saw a surge in critical reviews in 2012, which aligns with the introduction of the Acceptable Ads feature. These broader trends highlight the importance of long-term user feedback analysis. While immediate feedback post-update is crucial, observing larger sentiment trends over time can shed light on the long-term effects of various decisions.

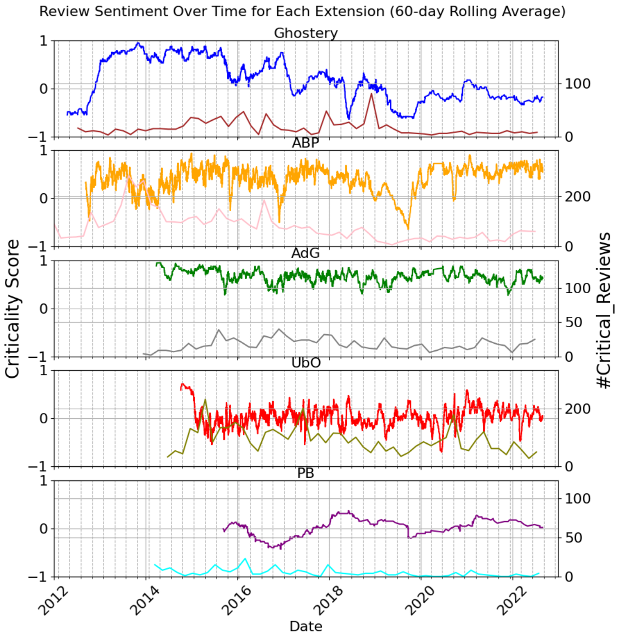

**Figure 5: Sentiment Vs Time plots for five extensions over 60-day rolling average. There seem to be a few noticeable trends but not a concrete pattern. For example, there seems to be a general decline in the sentiment for Ghostery from 2016 to 2020 and a steep rise in critical reviews for ABP around 2012.**

Variability in the criticality score around updates can indicate potential issues with new releases. Crowdsourcing has proven effective, especially in testing filterlists, yet the challenge remains in collecting enough feedback to evaluate each update accurately. To improve this, extensions should provide users with effective reporting tools.

# B MANIFEST V3 UPDATES

Manifest V3 governs Chrome plugin and extension API rules. New extensions must adhere to MV3 while existing ones can still use MV2, but Chrome ceased accepting new MV2 extensions in June 2022. According to Chrome, the motivation behind introducing MV3 was to improve extension performance and enhance user privacy. We find that most extensions do not have an overall adverse impact on performance. They do not significantly affect Load times and Data usage, and nearly all of them improve CPU performance. While numerous extensions exhibit subpar RAM usage when compared to the control case, it remains uncertain whether they fare better under MV3 standards.

# C WEBSITE CATEGORIZATION

Figure 6 shows the performance metrics for each extension, along with their corresponding tail values. We plotted the difference between each metric value in the extension and the control case over the inner page testing pool. Each subplot includes a median value

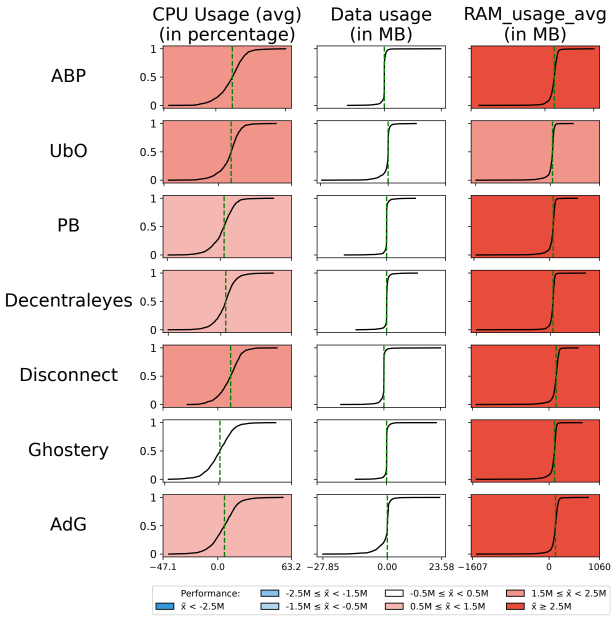

**Figure 6: Different performance metrics – CPU usage, Data usage, and RAM usage against each extension. The reported values are differences between the value recorded with the extension and the control case over the inner page testing pool. The *dotted* line represents the median of the data.**

line, which is highlighted based on the condition determined by the Median Absolute Deviation (M). MAD is a robust measure of variability for a univariate sample of quantitative data. Subplots displayed in blue indicate negative median values, indicating improved performance compared to the control case. Red subplots represent positive median values, indicating a decline in performance compared to the control case. We studied websites belonging to which categories are densely present at the tail ends. We refer to website categories [25] provided by the Interactive Advertising Bureau (IAB). We find that the website categories significantly differ for different metrics. For example, websites belonging to the Reference Materials category show high improvement in CPU usage in the presence of ad-blocking extensions. Similarly, Business and Industrial category websites show maximum improvement in load time. This shows us that different sets of websites react differently to various extensions based on the nature of the content they host. A further detailed analysis has been left for future work.

## D TOPIC FRAMEWORK

The privacy and usability topic framework can be found in Table 5. This framework serves as the foundation for user concerns and measurement methodologies mentioned in this paper. To enhance understanding, four additional columns have been included. The "#Reviews" column indicates the number of reviews falling under each broad category, while also reflecting the dataset size used for unsupervised LDA analysis. In the bracket, we represent the proportion of reviews within the respective general category out of the total number of critical reviews. The "Description" column provides a high-level definition for each broad category. Lastly, the last two columns, labeled "Non-Critical Review" and "Critical Review," present examples of respective categories of reviews, thus also highlighting the distinction between negative and critical reviews as per our proposed definition.

**Table 5: Privacy and Usability Topic Framework with supporting columns. #Reviews shows the number of reviews containing any one of the Related Keywords (the percentage of total critical reviews is represented in the bracket). Non-Critical and Critical Reviews columns reflect upon our definition of what type of reviews we consider as Non-Critical and Critical. The percentages represent the proportion of reviews within the respective general category out of the total number of critical reviews.**

| Broad Category | Related Keywords | #Reviews | Description | Non-Critical Review | Critical Review |
|---|---|---|---|---|---|
| block | block, prevent, protect, secure, detect, bypass | 4024 (32.01%) | This category covers reviews that talk about blocking and detecting ads/malware, preventing websites from getting rendered, etc. | it blocked the ads well | been using abp for a while, but seems to be blocking less ads over time. |
| ads | popup, pop-up, malvertising, cdn, content delivery, spyware, adware, malware, paid, tabs, notification, annoy | 911 (7.25%) | This category specifically covers popups, spywares, CDNs and other web objects that annoy people | this app is strong. i Love the interface. specially for using Block All Popups Button easily | Great idea, until ABP started throwing its own popups on my screen to ask me how it's going like a needy toddler. Just working against the problem. |
| break | break, broke, rendering, error, bug, stop, access, reload, load, unusable, crash, mess, cannot open, hang, fail, corrupt, fix | 3896 (30.99%) | This category covers site breakages, crashes, hangs, errors, etc. | Haven't had any ads bug and irritate me since installing this extension, Great service | Breaks a popular video website. 2/5. Do not recommend. |
| tracking | 3rd, third, party, tracking, icons, JavaScript, private, redirect, vpn | 315 (2.51%) | This category involves JavaScript and 3rd party tracking users and stealing private information | Just found this extension, no more probs against inline JavaScript ads. | It also blocks a variety of non-advertising related JavaScript, leading to rendering issues on multiple websites. |
| manual | manual, off, features, turn off, disable, click | 1326 (10.55%) | This category focuses on settings that users have to set manually, often causing them additional discomfort | This is honestly the best adblocker iv used. allows you to manually block ads that weren't automatically blocked. | Adblock still showing facebook's ads (even after adding the filter manually) |
| filter | remove, filter, list, whitelist, blocklist, blacklist, maintain | 1236 (9.83%) | This category covers filterlists that most extensions use to block ads, popups, etc. | Filters are powerful tools. If it blocks stuff on a website where you don't want it to block anything, you can just whitelist the site. | Blocks legit sites all the time. Mostly with Peter Lowe's crazy list. Adblock and Adblock plus don't have this issue. |
| config | default, config, configuration, sync | 287 (2.28%) | This category involves default and other configurations that are set for extension's standard functionality | Works fine, should have option for Windows Phone by default. | Great app, but it lacks default blocking options. There are services I want to block everywhere and there are services I don't want to block at all (e.g. facebook). |

| **Broad Category** | **Related Keywords** | **#Reviews** | **Description** | **Non-Critical Review** | **Critical Review** |
|---|---|---|---|---|---|
| priv_policy | privacy, policy, consent, information, install, false, security, anonymity, monetize | 11 (0.09%) | This category focuses on privacy policies of the extensions and any policy violations identified by users | I treasure my privacy on the internet and I have the right to privacy, and AdGuard definitely allows my privacy to be top priority. | works great normally, but for some reason chrome just. disables it without my consent? ive had to uninstall and reinstall multiple times because of this |
| compatibility | version, corrupt, browser, compatible, extension, chrome, firefox, disable | 2381 (18.94%) | This category covers compatibility issues around different browsers, user-agents, etc. | Works great! Allows me to use Microsoft Outlook Web Access (via Exchange 2010) in Ubuntu Linux with Chrome browser. | Very limited on options of browsers and devices. Go with something else unless they update the list. Only one flavor of Android, etc. |
| data | permission, data, encrypt, history, memory, storage, leak, sell, prefetch | 274 (2.18%) | This category deals with the handling of data and permissions | Works well, I like how it learns which trackers are bad or good and adjusts accordingly to keep your browser and personal data safe! | It whitelists every single ad company tracker. You have to manually set everything which means those crappy ad companies already got some of your data. |
| performance | efficient, inefficient, fast, light, slow, heavy, speed, memory, cpu, long, lag, delay, cost, ram | 756 (6.01%) | This category highlights performance gains and losses post extension installation | Best adblocker with less ram and cpu usage | It blocks the google cloud console and sometimes makes other websites slow. |