

Digital Disparities: A Comparative Web Measurement Study Across Economic Boundaries

Masudul Hasan Masud Bhuiyan

masudul.bhuiyan@cispa.de

CISPA Helmholtz Center for Information Security
Saarbrücken, Germany

Cristian-Alexandru Staicu

staicu@cispa.de

CISPA Helmholtz Center for Information Security
Saarbrücken, Germany

Matteo Varvello

matteo.varvello@nokia.com

Nokia Bell Labs
Holmdel, NJ, USA

Yasir Zaki

yasir.zaki@nyu.edu

New York University Abu Dhabi
Abu Dhabi, UAE

Abstract

While internet usage is slowly catching up globally, it is still unclear how the web experience differs in developing and developed countries. On the one hand, the web has a notoriously large inertia, with many webpages still relying on unencrypted HTTP, deprecated web features, or old and buggy libraries. On the other hand, developing countries are expected to leapfrog and directly adopt the newest technologies by learning from the prior mistakes of more developed countries. Anecdotal evidence suggests that webpages in developing and developed regions differ significantly. In this work, we test this hypothesis by measuring differences in web development practices across the two groups of countries, using multiple dimensions: webpages' size, complexity, security, privacy, quality, technology adoption, and accessibility. Concretely, we collect the largest dataset to date that compares web development practices across developed and developing regions – 200,000 webpages across 20 countries – which we aim to open source along with this publication. Our findings reveal that webpages in developing regions are generally smaller and simpler, utilizing fewer requests – an adaptation that improves the performance over slower network conditions common in these areas. However, these sites are less optimized in other critical aspects: they frequently employ inefficient image formats, include unnecessary JavaScript or CSS code, or lack responsive image designs. Notably, our security assessment shows developing regions lagging in HTTPS adoption and vulnerability mitigation, possibly due to lower awareness of best practices.

CCS Concepts

• **Social and professional topics** → **Computing / technology policy**; • **Information systems** → **World Wide Web**.

Keywords

Digital disparities, Web measurement, Web development practices

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW '25, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1274-6/25/04

<https://doi.org/10.1145/3696410.3714647>

ACM Reference Format:

Masudul Hasan Masud Bhuiyan, Matteo Varvello, Cristian-Alexandru Staicu, and Yasir Zaki. 2025. Digital Disparities: A Comparative Web Measurement Study Across Economic Boundaries. In *Proceedings of the ACM Web Conference 2025 (WWW '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3696410.3714647>

1 Introduction

According to the most recent 2023 data from International Telecommunication Union [7] and the World Bank [8], 67% of the global population uses the internet, with a 5% growth in 2023 alone. However, there are large discrepancies in internet adoption based on the economic situation of a country, *i.e.*, while 93% of the developed world is online, only 60% of the developing world is. Moreover, due to the limited availability of fixed broadband services and the prohibitive costs associated with them, users from developing countries mostly rely on mobile internet [5, 20]. This measurable inequality – also called digital disconnect [54] or digital divide [64] – is a missed opportunity for economic growth in developing countries [47].

It is commonly believed that digital inequality also extends to the web. For example, a recent work by Ruth et al. [74] finds that most websites visited in a country are specific to that country, with only few popular global websites. Moreover, the World Bank [47] notes that developed countries tend to host more web-based services. On the contrary, the lack of broadband [29, 48] and modern devices in developing countries might lead to web development optimizing for low-bandwidth conditions, *e.g.*, with more simplistic designs and supported features. Still, websites in these regions are often criticized for poor optimization, causing slow load times and high data usage [33, 79]. However, developing countries can leapfrog [32, 44] and adopt the latest technologies, such as WebAssembly, without the ossification observed for older webpages [58]. Also, users, developers, and regulators in developing countries might be insufficiently concerned with security and privacy.

To the best of our knowledge, no large-scale empirical study has systematically analyzed web development practices across different economic contexts [55]. Most existing research focus on macro-level indicators such as internet penetration rates, smartphone adoption, or general technological infrastructure [67, 77]. To address this research gap, we conduct a large-scale analysis comparing webpages from both developed and developing countries. Our study involves identifying and crawling 200,000 popular webpages, *i.e.*, 10,000 from the top 10 most populated developed and developing

nations, using tools like Google Lighthouse for performance audits and Puppeteer to capture detailed network data. This analysis considers multiple web metrics, which we use to explore websites' composition and performance, security, privacy, and technology adoption, as summarized below:

Webpage Size and Complexity: We assessed webpage size, HTTP requests, and multimedia usage to evaluate load times and data consumption. Websites in developing countries are generally smaller and simpler, thus more suitable for slower networks. Inefficiencies in image formats and JavaScript persist, highlighting opportunities for further optimization despite their reduced size.

Performance Optimization: We analyzed websites' performance by assessing the use of optimized image formats or responsive images or the inclusion of unnecessary JavaScript and CSS files. Despite smaller sizes, websites in developing countries waste more bandwidth due to poorly optimized images and unused code, lacking advanced optimizations, which might hinder user experience.

Security Measures: We evaluated security practices by examining HTTPS usage, outdated libraries, and CSP implementation. Developed countries have adopted HTTPS more widely, but surprisingly, they also contain a higher number of severe security vulnerabilities. Both groups show a widespread lack of CSP policies, leaving websites vulnerable to threats like cross-site scripting.

Privacy Practices: We analyzed tracking scripts and third-party cookies to assess privacy concerns. Websites in developed countries use more trackers and cookies, likely due to data-driven business models. Despite stricter privacy regulations, tracking remains widespread via consent mechanisms. Developing countries have fewer trackers, but weaker privacy regulations remain a concern [25].

Technology Adoption: We examined the adoption of modern web APIs and protocols to assess if developing countries leapfrog older technologies. Surprisingly, some developing countries adopt HTTP/2 and HTTP/3 at rates comparable to or higher than developed ones, likely due to fewer legacy systems. Both groups, however, show similar web API adoption.

2 Related Work

Multiple studies have quantified cellular and broadband connectivity in developing regions, performing measurements, and proposing solutions to improve connectivity. Hassan et al. [51–53] investigated cellular connectivity in rural countries and how solutions such as Software Defined Networking can provide affordable connectivity. Similarly, Marshini et al. [38] explored mobile broadband connectivity in South Africa, while Zaki et al. [78] investigated connectivity challenges in Ghana. Fanou et al. [41] further highlighted the disparities in web content infrastructure across Africa. Research by Koradia et al. [59], Sharma et al. [75], and Bischof et al. [31] focused on cellular network issues in India and Cuba.

Google's AMP [45] and Facebook Lite [6] are products designed to address web complexity in developing regions. AMP provides a framework for faster web page construction, while Facebook Lite is optimized for low-end phones and slow networks. Google also launched Web Light [13], which transcoded webpages to lighter versions for slow connections. Some research papers also developed concrete web optimizations: Lite-Web [34], for example, aims

to reduce web complexity through tools such as SlimWeb [35], JSCleaner [36], and Muzeel [61].

Web performance in developing regions are worsened by the diffusion of low-end devices. Naseer et al. [65, 66] noted that JavaScript execution often triggers memory constraints, while rising web complexity and JavaScript-based tracking inflate costs for users. Amjad et al. [27] showed that selective JavaScript blocking can help, though with challenges. Additional research has tackled web accessibility in these areas. Raza et al. [71] proposed enhanced caching to compensate for limited cloud infrastructure, while Ahmad et al. [26] advocated for a new content ecosystem tailored to local needs, essential for users in resource-constrained environments.

Finally, Habib et al. [50] introduced a fairness metric to redesign the web with a focus on affordability and inclusivity. Unlike [49, 50], we analyze and explore 200k popular web pages, from the most populated developed and developing regions, in terms of their performance, security, privacy, and technology adoption. Additionally, in comparison to [74], which focuses on web traffic distribution and the most visited categories, we delve into the finer details of these pages, comparing factors such as image formats, unnecessary JavaScript usage, and security measures.

3 Methodology

Our methodology combines Google Chrome User Experience (CrUX) dataset [14] with the International Monetary Fund (IMF)'s classification, and with web data (webpage composition, protocol usage, etc.) we crawled using Google Lighthouse and Puppeteer. IMF classifies countries as "advanced" and "emerging market and developing economies" based on per capita income, export diversification, and financial system integration. CrUX data provides real-world web performance metrics like page load times and responsiveness, offering insights into website performance across regions.

3.1 Country Selection Criteria. The countries in this study are selected to ensure a balanced representation of both developing and developed economies, following the classifications provided by the IMF. Specifically, we select a sample of 10 developing and 10 developed countries for a comparative analysis of their web ecosystems. The selection process was guided by two primary criteria: population size and the availability of 10,000 websites per country that met our inclusion criteria (§3.2) within the CrUX database [14].

For developing countries, we identified the ten most populous nations from a cohort of 152 economies classified as "developing" by the IMF [24]. Population size is chosen as the primary criterion to ensure that the selected countries represent a significant portion of the global population within the developing world. Ethiopia, although the 10th most populous developing country [12], was excluded because we could not identify at least 10,000 websites that met our criteria (§3.2). To complete our sample, we included the Philippines in place of Ethiopia. The ten developing countries selected, in order of population, are: China, India, Indonesia, Pakistan, Brazil, Nigeria, Bangladesh, Russia, Mexico, and the Philippines.

For the developed countries, we focused on the most populous nations classified as "Advanced Economies" by the IMF. This selection ensured that major global economies were adequately represented in our dataset. The ten developed countries, ordered by population,




Figure 1: Comparison of website size (a) and request count (b) between developed and developing countries.

are: the United States, Japan, Germany, France, the United Kingdom, Italy, South Korea, Spain, Canada, and Australia.

Together, the developing countries in our dataset account for a combined population of 4.3 billion people, representing 54.33% of the global population [23]. In comparison, the developed countries in our sample account for 906.01 million people, or 11.3% of the world’s population. By using population size and CrUX data availability as selection criteria, our study captures a significant portion of the world’s population across both economic classifications, providing a robust foundation for comparing the web ecosystems of developing and developed nations.

3.2 Website Selection. To investigate the digital landscape in each country, we retrieve data on 10,000 websites per country from the CrUX dataset. A website is categorized under a particular country if it either has a country code top-level domain (ccTLD) for that country (e.g., de for Germany) or if its Whois [40] data indicates a physical registration address in that country. This criterion was used to mitigate the over-representation of globally dominant websites, which could skew the analysis of local web ecosystems. We utilized the `python-whois` package and the `who.is` API [22] to retrieve registration information for the websites. We start from the 10,000 most popular websites for each country – as indicated by the CrUX list – and then apply our selection criteria to filter the relevant websites for analysis. This approach ensures a uniform number of websites for each country, and a comparable distribution of website rankings. In most cases, our 10,000 websites were within the top 50,000 of that country. However, for Bangladesh, Pakistan, Nigeria, and the Philippines, we had to include websites that were less popular to meet our criteria and target sample size. Please refer to Appendix A.1 for the distribution of website rankings.

3.3 Crawler Configuration. We crawl the selected websites using Google Lighthouse [11] for performance audits, while simultaneously employing Puppeteer [15] to intercept and log network requests and responses, capturing detailed data such as request headers, responses, and metadata. In the presence of unreachable websites, e.g., due to temporary outages, Lighthouse timeouts, certificate errors, or other technical problems, we supplemented the list with lower-ranked (less popular) websites. For the Lighthouse audits, we emulated an iPhone X mobile environment, with a screen resolution of 375x812 pixels and a device scale factor of 3, ensuring that the mobile version of each page was loaded. This configuration was chosen to reflect common mobile browsing conditions and to avoid bias due to default client configuration as observed in [57]. We run our crawlers across three server machines hosted at two university campuses (Germany and UAE). This selection was

motivated by the minimal crawling bias observed for university IP addresses when compared to, for instance, cloud IP addresses [57].

3.4 Library Detection. To analyze the usage of JavaScript libraries across the selected websites, we employ multiple methods for library detection. We use the `retirejs` Python library [17], custom regular expressions to extract library information from request URLs, and data from Google Lighthouse audits. This approach helps capture a broader range of library usage across websites.

3.5 Limitations. Our crawler operates from Germany and UAE for scalability, though Germany’s inclusion as a test country may introduce bias. Ideally, crawlers would run within each studied country, but logistical constraints made this unfeasible. VPNs proved unreliable, as VPN detection could skew results, and many rely on cloud servers that lack tracking domains and JavaScript libraries [57]. Since crawlers don’t operate within most studied countries or simulate realistic mobile networks, they cannot fully capture Web performance tied to user experience. To mitigate this, we use CrUX data, offering real-world metrics from actual devices and networks.

Google Lighthouse also has limitations. By default, it doesn’t handle cookie banners, which may prevent full webpage rendering. While solutions exist [16], their reliability is untested. Lighthouse also cannot handle user authentication, affecting performance and security metrics [70]. Additionally, `retirejs` may miss libraries, especially in bundled files, limiting library detection accuracy [69].

4 Size and Complexity

4.1 Static Footprint

We measure a webpage size as the total number of bytes required to load all its content. Figure 1(a) shows that the median website size in developed countries is 3.35 MB, *i.e.*, slightly larger than the 2.99 MB median size in developing countries, both of which exceed the global median of 2.6 MB as per the HTTP archive [28]. Similarly, websites in developed countries require a median of 108 requests to load (Figure 1(b)), compared to 89 requests in developing countries – while the global median is 70 requests. These differences between developed and developing countries are statistically significant ($p < 0.001$), highlighting the larger footprints of websites in developed countries. South Korea stands out with exceptionally large websites – median size of 5 MB and upper whisker close to 30 MB – due to higher multimedia content and interactive features supported by the country’s fast internet speeds [19].

This result contrasts with earlier findings that focused on public service websites [49], where websites in developing countries were significantly larger and less optimized. In our broader comparison of general websites, however, the opposite is true: websites in




Figure 2: Bandwidth allocation for different resource types.




Figure 3: The sum of all image sizes and the number of image requests across developed and developing regions.

developing countries have smaller footprints and fewer resource requests, indicating they are more suitable for slower networks. Despite this, both developed and developing countries exceed the global median in terms of website size and number of requests, suggesting further optimization could reduce data costs, especially for users with limited access to affordable internet services.

4.2 Breakdown by Resource Type

To understand which webpage characteristics contribute to the larger sizes of websites in developed countries, we analyze the sizes of different resources across pages. We distinguish between HTML documents, images, JavaScript, and third-party content, *i.e.*, any resource imported from a third party domain [80]. Figure 2 shows a clear contrast in how websites from these two regions allocate bandwidth across webpage components, particularly images. Websites in developing countries dedicate a much larger share of their page size to images, ranging from 48% to 57%. For instance, websites in Indonesia allocate 55.3% of their page size to images, compared to websites in developed countries like the US (28.8%), the UK (36.3%), and Canada (37.2%), where image bandwidth usage is < 40%.

Figure 3 further focuses on image data usage, revealing that websites in developing countries devote a larger proportion of their bandwidth to images, even though they make fewer image requests (19 on average) than websites in developed countries (24 on average). Notably, websites in developing regions allocate more than 50% of their bandwidth to images, averaging 822.75 KB, compared to websites in developed regions, which allocate less than 40% and average 839.62 KB. This indicates that the higher share of image-related bandwidth in developing countries stems from larger image files rather than a larger number of image requests.

By analyzing image size distribution, we find that, indeed, websites in developing countries use significantly larger images compared to those in developed countries (see Appendix A.2). For example, 15% of images on websites in developing countries exceed 100 KB, whereas this is the case for only 12% of images in developed countries. This reinforces the observation that websites in developing countries do not host more images but rather use larger images, which increases overall website size and bandwidth consumption. This leads to potential negative effects, particularly for users in regions with slower internet speeds and limited data plans.

4.3 Dynamic Footprint

We here focus on JavaScript (JS), the dominant factor in a website’s performance [37], with Chrome spending over 30% of its processing time on it [3]. Figure 4(a) shows the difference in the total size of JS usage across various countries. The median JS size for websites in developed countries such as the US and Australia is 921.71 KB and 968.61 KB, respectively, while in developing countries like Bangladesh and Pakistan, the median JS size is significantly lower, at 466 KB and 547 KB. This trend is consistent across regions, with developed countries generally using more JavaScript code. The overall median JS size for developed countries is 725 KB, compared to 546 KB for developing countries. This shows a clear difference in JS usage, suggesting that websites in developed countries rely more on JS to create richer, more interactive user experiences.

Further supporting this trend, Figure 4(b) shows that websites in developed countries make more JS network requests than those in developing regions. For example, the median number of requests for websites in the US and Australia is 25 and 29, respectively, while in developing nations like Bangladesh and Pakistan, the median number of requests is 21 and 25, respectively. Overall, the median number of JS requests is 21 for developed and 18 for developing countries. This difference highlights again the greater reliance on JS in developed countries. When combined with the median document size (see Appendix A.3), which includes all HTML files, where developed countries have a median size of 33.6 KB compared to 26.1 KB in developing regions. It becomes evident that the increased JS usage contributes to the overall websites’ size difference.

Figure 5 shows that this complexity is also reflected in the DOM size, or number of DOM elements in a webpage, which represents the structure and complexity of the page. A larger DOM size generally indicates a more complex webpage with more interactive elements, often driven by extensive JS usage. For example, the median number of DOM elements in developed countries like Germany (1,283) and Italy (1,058) is significantly higher than in developing countries such as Pakistan (940) and Russia (926). The overall median DOM size for developed countries is 969 elements, while in developing countries, it is much lower at 633 elements. This larger DOM size in developed countries highlights the increased complexity and interactivity of websites, which is often due to the heavier use of JS to enhance dynamic content and features.

The higher complexity of webpages in developed countries, as indicated by larger document sizes and DOM elements, also impacts browser performance, particularly in terms of *total blocking time (TBT)*, a metric that measures how long a browser’s main thread is blocked by tasks, preventing users from interacting with the page. Websites in developed countries have a significantly larger TBT, with median values of 631.0 ms compared to 418.5 ms in developing countries (see Appendix A.4). This longer TBT is a direct consequence of larger amounts of JS and HTML code, which must be interpreted before the website becomes fully interactive.

5 Performance Optimizations

5.1 Image Compression

Optimizing image formats and using properly sized images are critical for improving web performance, particularly in regions with bandwidth constraints. Modern optimized formats such as




Figure 4: Comparison of total JS size (a) and number of JS requests (b) between developed and developing countries.




Figure 5: Total number of DOM elements across countries.

WebP, AVIF, and SVG are recommended by best practices due to their superior compression capabilities and ability to maintain high image quality with smaller file sizes [18]. While JPEG 2000 (JP2) is not recommended, studies have shown that this format outperforms traditional JPEG in terms of compression efficiency [9].

Figure 6(a) shows the analysis of image formats across countries, which reveals disparities in the adoption of these optimized formats. Developed countries have a higher usage of modern formats like WebP, AVIF, and SVG, contributing to faster page loads and more efficient data usage. For instance, in the United States, modern formats account for approximately 37% of image usage, with SVG at 20.03%, WebP at 13.50%, and AVIF at 3.81%. Traditional formats like JPEG and PNG account for 17.90% and 24.52%, respectively. Similarly, Germany shows substantial adoption of modern formats, with SVG at 22.45%, WebP at 11.33%, and AVIF at 2.69%, totaling about 36% of modern format usage. In contrast, developing countries have lower adoption rates of modern formats; in India, modern formats constitute approximately 28% of image usage—SVG at 15.91%, WebP at 9.64%, and AVIF at 2.68%—while traditional formats like JPEG and PNG dominate at 23.24% and 29.62% respectively. Nigeria exhibits even lower modern format usage at around 24%, with SVG at 14.95%, WebP at 9.02%, and AVIF at 0.52%, whereas JPEG and PNG account for 27.48% and 31.13%.

Next, we analyze image compression wastage, which refers to the proportion of image file sizes that exceed what is necessary for maintaining visual quality. Lighthouse computes this by comparing the size of the rendered image against the actual image size, accounting for the device’s pixel ratio to assess whether images are optimally sized. Figure 6(b) shows that developed countries generally have lower image compression wastage rates compared to developing countries. Developed countries like the United States, South Korea, and Canada, which have median wastage rates of approximately 38-39%, also exhibit higher usage of modern formats such as WebP and AVIF in Figure 6, contributing to better management of image sizes and reducing wastage. In contrast, developing countries like India and Bangladesh, with higher wastage rates of 55% and 75%, respectively, show a lower adoption of optimized

formats, indicating that inefficient image formats and large file sizes are more commonly used, particularly on smaller devices. This correlation between format usage and wastage emphasizes the importance of adopting optimized formats to improve performance.


5.2 Unused JavaScript and CSS Code

The unused JavaScript (JS) audit highlights code that is downloaded and parsed but never executed, contributing to significant performance inefficiencies. Although this code isn’t directly used during page load, it still incurs costs as browsers must download, parse, and sometimes compile it. While tools like Lighthouse help identify this issue, they may overestimate the amount of unused JS by not fully simulating user interactions that might trigger the code later.


Several factors contribute to the presence of unused JS, including the use of pre-built libraries, legacy code, and a lack of regular optimization. Developers often include entire libraries or modules even when only small portions are needed. Additionally, outdated or redundant code can accumulate on websites, remaining in the assets due to limited maintenance or the complexity of refactoring. The increasing reliance on large JS frameworks exacerbates this issue, as these frameworks are often not tailored to the specific needs of individual webpages. Unused JS significantly impacts performance, especially on mobile devices. Mobile phones, particularly low-end models common in developing countries, require up to seven times more processing power than desktop devices to handle JS [37], further slowing down page loads. In regions with slower networks and higher data costs, the burden of downloading and processing unnecessary code is even more problematic.

On average, 55% of JS is unused across all countries, as shown in Figure 7. Developed nations like the United States (82.65%) and Japan (79.76%) show high inefficiencies, while countries like South Korea (93.69%) and Germany (92.65%) fare even worse. The problem is most severe in developing nations such as Bangladesh (99.68%) and China (98.72%), where limited network speeds and strained resources amplify the negative impact of unused JS. CSS wastage is also prevalent, with over 91% of global CSS rules remaining unused. Developing nations like Pakistan and Bangladesh exhibit wastage levels above 95%, while developed countries like Japan and Germany report slightly lower rates, around 92%. This wastage is likely due to outdated frameworks and poor mobile optimization, as discussed further in Appendix A.6.

It is crucial to address these inefficiencies, especially in developing regions where slow networks and low-end devices exacerbate the issue. Reducing unused JS and CSS through better development practices will improve performance, particularly for users in regions with limited network and device capabilities.



(a) Distribution of image format by country.



(b) Wasted image bytes (%) from compression inefficiencies by country.

Figure 6: Comparison of image format distribution (a) and percentage of wasted bytes due to compression inefficiencies (b) across different countries.




Figure 7: Distribution of unused JavaScript across countries.




Figure 8: Distribution across developing and developed regions of the percentage of mobile webpage loads which were Good, Needs Improvement, or Poor.

5.3 Webpage Performance

We conclude the section by analyzing the quality of browsing experience for users in both developed and developing regions. Using the CrUX dataset [14], which anonymously reports web performance metrics from mobile users, we focus on the Largest Contentful Paint (LCP) [10]. LCP measures how quickly the main content of a webpage becomes visible, a key metric recommended by Google as a reliable indicator of page load performance [10]. LCP performance is classified into three categories: Good (≤ 2.5 sec), Needs Improvement (2.5-4 sec), and Poor (≥ 4 sec) [2]. CrUX reports the percentage of mobile visits in each country that fall into these categories for the pages under test. Although CrUX does not disclose the total number of reports for privacy reasons, we estimate a minimum of 70.2 million reports [43]. This provides valuable insights into real-world user experiences across different regions.


For visibility reasons, Figure 8 aggregates the results across developing and developed regions, and shows the percentage of mobile visits which were Good, Needs Improvement, or Poor. For example, 50% of the developed webpages have close to 60% of their mobile visits, which were “Good”, whereas this number drops to 50% for developed regions. Conversely, developing regions have overall more webpages characterized by either “Needs Improvement” or “Poor” performance. This analysis confirms that users in developing regions experience a slower web, however the impact is less

dramatic than expected. This is likely due to the overall smaller sizes, and lack of complex JS functionalities, previously discussed.


6 Privacy & Security Analysis

Trackers and Third-party Cookies. They play a crucial role in how websites collect data on user behavior, preferences, and interactions. These tools enable personalized advertising, content customization, and detailed analytics. Trackers work by embedding scripts in websites to monitor user actions, while third-party cookies allow tracking across multiple domains. However, these practices raise privacy concerns, as they allow companies to build detailed user profiles. Surprisingly, despite stricter regulations like GDPR, our analysis, as shown in Figure 9(a), indicates a counter-intuitive trend: websites in developed countries tend to use more trackers despite the presence of these regulations. For example, 70.86% of websites in Germany have trackers, while this percentage increases to 93.07% in the United Kingdom and 92.06% in Australia. Other developed countries, such as Canada (89.50%), France (83.96%), and Japan (90.41%), also show high tracker prevalence. On average, websites in developed countries employ approximately nine trackers per site. In contrast, the prevalence of trackers in developing countries shows more variability. In China, 68.01% of websites have trackers, while Bangladesh (90.77%) and Pakistan (90.26%) exhibit similarly high tracker usage. Countries such as Brazil (89.55%) and Russia (91.52%) also have widespread tracker deployment. On average, websites in developing countries use six trackers, significantly fewer than their counterparts in developed countries. The higher tracker prevalence in developed nations can be explained by the fact that these countries have more mature advertising industries and e-commerce platforms that rely heavily on user data for targeted marketing and sales optimization. Additionally, data-driven business models are more prevalent, making user data a valuable asset.

In Figure 9(b), we show the distribution of third-party cookies across both developed and developing countries, showing a low overall prevalence, likely driven by changes in browser policies, e.g., Google Chrome is planning to soon drop support for such cookies. However, Russian websites stand out with a particularly high prevalence of both trackers and third-party cookies. The median number of trackers on Russian websites is 17, and 86% of these websites employ third-party cookies. This extensive use of third-party cookies and trackers in Russia is likely driven by a strong reliance on data-driven advertising models and lack of strict privacy-related legislation. These results are consistent with those of Gotze et al. [46],



(a) Prevalence of trackers across countries



(b) Percentage of websites using third-party cookies across countries

Figure 9: Comparison of the prevalence of trackers and third-party cookies across developed and developing countries.




Figure 10: HTTPS adoption across regions.




Figure 11: CSP usage rates across regions.




Figure 12: Number of vulnerabilities by country.

who also show that almost 90% of Russian websites extensively use both trackers and third-party cookies.

HTTPS Adoption. HTTPS [72] encrypts web content protecting it from interception and manipulation. Figure 10 shows the adoption of HTTPS across both developed and developing nations. In developed countries, only 5% of websites still use unencrypted HTTP, a significant improvement from previously-reported measurements [42]. In contrast, 12% of websites in developing countries still rely on HTTP. This gap is even more concerning in certain countries, with Bangladesh and China having the most websites using HTTP: 25% and 20%, respectively (see Appendix A.5). These results show a lag in adopting basic security standards in the developing world, potentially on purpose, leaving users in those regions vulnerable to (un)lawful interception.

Content Security Policy (CSP) Usage. Content Security Policies (CSPs) act as a critical defense mechanism against cross-site scripting (XSS) and other injection attacks, offering a structured way to control which resources can be loaded and executed on a webpage. However, the correct usage of CSP is a known challenge [73]. Figure 11 illustrates CSP usage rates across countries, showing a significant lack of adoption in both developed and developing nations. In developed countries, 75% of websites fail to implement a CSP. The situation is even more concerning in developing countries, where 81.5% of websites do not implement CSP at all. This widespread absence of valid and robust CSPs is not confined to specific regions; over 99% of websites globally, in both developed and developing countries, either lack CSPs altogether or use highly ineffective deployments. These findings underscores a universal struggle with content security policies, as previously reported by developers in the recent work of Roth et al. [73].


Vulnerabilities in Web Ecosystems. When analyzing vulnerabilities, an unexpected trend emerges: websites in developed countries are more prone to severe security issues than those in developing regions. Figure 12 provides a comparison of vulnerability rates across countries, highlighting the higher incidence of critical




Figure 13: Outdated libraries usage across countries.

vulnerabilities in developed nations. For example, in developed countries such as Italy, Canada, Spain, Australia, and South Korea, around 10% of websites exhibit high or critical vulnerabilities. In contrast, the highest rate in the developing world is 7%, observed in Bangladesh. This trend is particularly interesting because, despite higher security standards, developed countries experience more critical security issues. One plausible explanation for this trend is the increased reliance on JS libraries and frameworks in developed regions (see Figure 2). While JS enhances website functionality, it also increases the attack surface. The vast ecosystem of third-party libraries and dependencies, which are not always well-maintained, can contribute to the large amount of observed vulnerabilities. These results are in line with prior measurements [62, 63] that report that 40% of the websites contain a client-side vulnerability.

Outdated and Duplicate Libraries. The use of outdated and/or duplicate libraries can significantly increase a website’s exposure to security risks, as these libraries may harbor vulnerabilities that attackers can exploit. Figure 13 illustrates the prevalence of outdated libraries across different countries. Developers across regions often hesitate to upgrade to newer versions of libraries. These results are in line with Raula et al. [60] who state that 81.5% of systems include at least one outdated library, and Lim et al. [63] who say that the majority of websites include JavaScript libraries that are more than seven years old. The issue is especially critical in Korean websites, where 94% of websites use at least one older library version, contributing to the higher number of vulnerabilities observed in Korean websites, as shown in Figure 12.



(a) Adoption of HTTP protocols across countries



(b) Usage of web APIs across countries

Figure 14: Comparison of the adoption of HTTP protocols and web APIs across developed and developing countries.

7 Technology & Affordability

Technology Adoption. The adoption of new technologies reflects factors like economic development and infrastructure. While developed countries are assumed to adopt new technologies faster, data on HTTP protocols challenges this notion.

Figure 14(a) shows HTTP protocol usage across countries. Newer protocols like HTTP/2 (h2) and HTTP/3 (h3) improve performance and security. Surprisingly, websites in developing countries are not more reliant on older protocols than developed nations. Nigerian websites, for example, show high combined h2 and h3 usage at 65% (37.80% for h2, 27.09% for h3), exceeding the US at 58% (31.87% h2, 26.17% h3). Similarly, Pakistani websites report strong adoption with 36.03% for h2 and 29.39% for h3. In contrast, some developed countries like Germany and the UK report lower combined adoption of h2 and h3. Outdated protocols like HTTP/1.0 remain rare but slightly more common in developed countries (e.g., US: 0.31%, Nigeria: 0.02%). These findings challenge assumptions about technological leadership.

Figure 14(b) shows the adoption of various Web APIs; we first selected 74 Web API standards using the same technique as [76], and then focused on 10 APIs that are security- and privacy-critical. The results show that the adoption of these security-critical APIs is not confined to developed countries. Russian websites show significant use of the Web Cryptography API (8,219 instances vs. 4,214 in the US). Similarly, Bangladeshi websites surpass US sites in Encrypted Media Extensions usage (1,307 vs. 887 instances). Nigerian websites also demonstrate active engagement with modern web technologies, with 2,318 instances of the Web Cryptography API, highlighting strong adoption in developing regions. These findings indicate that developing countries are embracing new technologies, sometimes at rates matching or surpassing developed nations. This may be due to the absence of legacy systems, enabling easier adoption of modern technologies, alongside widespread mobile internet use and a younger, tech-savvy population.

Website Affordability. To assess websites affordability, we use the PAW (Price Adjusted Web access) index [68]. This index measures the reduction in webpage sizes needed to meet the UN Broadband Commission’s target, stating that broadband services should not exceed 2% of a country’s Gross National Income per capita [39]. A PAW index of 1 or less indicates affordable access, while values above 1 suggest that webpage sizes or broadband costs are too high relative to income, limiting access to essential online services [68]. See Appendix A.7 for details on PAW calculation.


**Figure 15: Website affordability (PAW Index) across countries.**

Figure 15 shows the PAW indices for 20 countries, grouped into developed and developing. The index values were based on average webpage sizes and mobile broadband prices in each country. The median PAW index for both developed and developing countries was below 1, indicating overall affordability. Japan (1.11) was the only developed country with a PAW index above 1. In developing countries, Pakistan (1.14), Nigeria (1.04), Bangladesh (1.19), and the Philippines (1.54) exceeded the threshold, with the Philippines having the highest index. This points to substantial barriers to web access, due to a combination of high mobile data costs and large webpages. Developed countries showed less variation in PAW values, while developing countries had more inconsistency in affordability across regions, reflecting significant regional disparities.

8 Conclusion

Despite two-thirds of the global population being online, a significant digital divide persists: 93% of people in developed regions are connected, compared to just 60% in developing ones. This gap, exacerbated by reliance on mobile internet and high data costs, affects access and user experience, with slower load times, and accessibility challenges (see Appendix B) in developing regions. This study examines these disparities by comparing the web in developed and developing regions. Using Google CrUX, we identified the 10,000 most popular websites in the top 10 countries from each group. Next, using Google Lighthouse and Puppeteer, we crawled and audited a total of 200,000 webpages. The resulting dataset, the largest of its kind, is publicly available on GitHub [4] (Appendix C). Our findings show websites in developing regions are 10% smaller and use fewer third-party resources and resource-heavy JavaScript, leading to simpler DOM structures and shorter browser blocking time. However, these sites exhibit weaker optimization, wasting 10% more bandwidth on improperly sized images and redundant JS/CSS. Their smaller size seems to stem from simpler structures rather than intentional optimization. Security practices also lag, with HTTPS adoption at 88% (vs. 95%) and CSP adoption at 75% (vs. 81.5%). Whether these gaps reflect limited awareness or design choices for surveillance remains unclear.

References

- [1] [n. d.]. axe-core/doc/rule-descriptions.md at develop · dequelabs/axe-core. <https://github.com/dequelabs/axe-core/blob/develop/doc/rule-descriptions.md>
- [2] [n. d.]. Core Web Vitals report - Search Console Help. <https://support.google.com/webmasters/answer/9205520?hl=en>
- [3] [n. d.]. The cost of JavaScript in 2019 - V8. <https://v8.dev/blog/cost-of-javascript-2019>
- [4] [n. d.]. Dataset for the paper "Digital Disparities: A Comparative Web Measurement Study". <https://github.com/kal-purush/digital-disparities-www25>
- [5] [n. d.]. Desktop vs Mobile vs Tablet Market Share in Asia. <https://gs.statcounter.com/platform-market-share/desktop-mobile-tablet/asia> Accessed: 2024-10-07.
- [6] [n. d.]. Facebook Lite APK for Android. <https://www.facebook.com/lite/>
- [7] [n. d.]. Individuals using the Internet - ITU DataHub. <https://datahub.itu.int/data/?e=1&i=11624>
- [8] [n. d.]. Individuals using the Internet (% of population) | Data Catalog. <https://datacatalog.worldbank.org/indicator/715ee78d-bfce-eb11-bacc-000d3a596ff0/Individuals-using-the-Internet---of-population->
- [9] [n. d.]. JPEG vs. JPEG 2000: Which is better? | Adobe. <https://www.adobe.com/creativecloud/file-types/image/comparison/jpeg-vs-jpeg-2000.html>
- [10] [n. d.]. Largest Contentful Paint (LCP) | Articles. <https://web.dev/articles/lcp>
- [11] [n. d.]. Lighthouse. <https://developer.chrome.com/docs/lighthouse>
- [12] [n. d.]. List of 152 Developing Countries. <https://www.worlddata.info/developing-countries.php> Accessed: 2025-01-29.
- [13] [n. d.]. Load mobile pages faster with Web Light - Google Search Help. <https://support.google.com/websearch/answer/9836344?hl=en>
- [14] [n. d.]. Overview of CrUX | Chrome UX Report. <https://developer.chrome.com/docs/crux>
- [15] [n. d.]. Puppeteer | Puppeteer. <https://pptr.dev/>
- [16] [n. d.]. Removing cookie consent banners from your Lighthouse test. <https://pagevitals.com/docs/guides/removing-cookie-consent-banners-from-your-lighthouse-tests/>
- [17] [n. d.]. retirejs: Port of RetireJS in Python. <https://github.com/FallibleInc/retirejslib>
- [18] [n. d.]. Serve images in modern formats | Lighthouse. <https://developer.chrome.com/docs/lighthouse/performance/uses-webp-images>
- [19] [n. d.]. Speedtest Global Index - Internet Speed around the world. <https://www.speedtest.net/global-index#mobile>
- [20] [n. d.]. The State of Mobile Internet Connectivity Report 2023 - Mobile for Development. <https://www.gsma.com/r/somic/> Accessed: 2024-10-07.
- [21] [n. d.]. Web Content Accessibility Guidelines (WCAG) 2.1. <https://www.w3.org/TR/WCAG21/>
- [22] [n. d.]. WHOIS Search, Domain Name, Website, and IP Tools - Who.is. <https://who.is/>
- [23] [n. d.]. World Bank Open Data. <https://data.worldbank.org>
- [24] [n. d.]. World Economic Outlook Database - Groups and Aggregates. <https://www.imf.org/en/Publications/WEO/weo-database/2023/April/groups-and-aggregates>
- [25] 2020. Data and privacy unprotected in one third of countries, despite progress | UNCTAD. <https://unctad.org/news/data-and-privacy-unprotected-one-third-countries-despite-progress>
- [26] Talal Ahmad, Yasir Zaki, Thomas Pötsh, Jay Chen, Arjuna Sathiaselalan, and Lakshminarayanan Subramanian. 2019. GAUIS: a new mobile content creation and diffusion ecosystem for emerging regions. In *Proceedings of the Tenth International Conference on Information and Communication Technologies and Development (Ahmedabad, India) (ICTD '19)*. Association for Computing Machinery, New York, NY, USA, Article 34, 5 pages. doi:10.1145/3287098.3287130
- [27] Abdul Haddi Amjad, Zubair Shafiq, and Muhammad Ali Gulzar. 2023. Blocking JavaScript without Breaking the Web. In *Privacy Enhancing Technologies Symposium (PETS)*.
- [28] HTTP Archive. 2021. State of the Web. <https://httparchive.org/reports/state-of-the-web>.
- [29] World Bank. 2016. *World Development Report 2016: Digital Dividends*. World Bank Publications. <https://www.worldbank.org/en/publication/wdr2016>
- [30] Masudul Hasan Masud Bhuiyan, Matteo Varvello, Cristian-Alexandru Staicu, and Yasir Zaki. 2025. Non-Western Perspectives on Web Inclusivity: A Study of Accessibility Practices in the Global South. In *Proceedings of the 47th International Conference on Software Engineering (ICSE 2025)* (Ottawa, Canada). IEEE / ACM.
- [31] Zachary S. Bischof, John P. Rula, and Fabián E. Bustamante. 2015. In and Out of Cuba: Characterizing Cuba's Connectivity. In *Proceedings of the 2015 Internet Measurement Conference* (Tokyo, Japan) (IMC '15). Association for Computing Machinery, New York, NY, USA, 487–493. doi:10.1145/2815675.2815718
- [32] Elise S Brezis, Paul R Krugman, and Daniel Tidson. 1991. Leapfrogging: A theory of cycles in national technological leadership.
- [33] Moumena Chaqfeh, Rohail Asim, Bedoor AlShehli, Muhammad Fareed Zaffar, Talal Rahwan, and Yasir Zaki. 2023. Towards a world wide web without digital inequality. *Proceedings of the National Academy of Sciences* 120, 3 (2023), e2212649120.
- [34] Moumena Chaqfeh, Rohail Asim, Bedoor AlShehli, Muhammad Fareed Zaffar, Talal Rahwan, and Yasir Zaki. 2023. Towards a world wide web without digital inequality. *Proceedings of the National Academy of Sciences* 120, 3 (2023), e2212649120.
- [35] Moumena Chaqfeh, Muhammad Haseeb, Waleed Hashmi, Patrick Inshuti, Manesha Ramesh, Matteo Varvello, Lakshminarayanan Subramanian, Fareed Zaffar, and Yasir Zaki. 2023. To Block or Not to Block: Accelerating Mobile Web Pages On-The-Fly Through JavaScript Classification. In *Proceedings of the 2022 International Conference on Information and Communication Technologies and Development* (Seattle, WA, USA) (ICTD '22). Association for Computing Machinery, New York, NY, USA, Article 25, 12 pages. doi:10.1145/3572334.3572397
- [36] Moumena Chaqfeh, Yasir Zaki, Jacinta Hu, and Lakshmi Subramanian. 2020. JSCleaner: De-Cluttering Mobile Webpages Through JavaScript Cleanup. In *Proceedings of the Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 763–773. doi:10.1145/3366423.3380157
- [37] Moumena Chaqfeh, Yasir Zaki, Jacinta Hu, and Lakshmi Subramanian. 2020. Jscleaner: De-cluttering mobile webpages through javascript cleanup. In *Proceedings of the Web Conference 2020*. 763–773.
- [38] Marshini Chetty, Srikanth Sundaresan, Sachit Muckaden, Nick Feamster, and Enrico Calandro. 2013. Measuring broadband performance in South Africa. In *Proceedings of the 4th Annual Symposium on Computing for Development* (Cape Town, South Africa) (ACM DEV-4 '13). Association for Computing Machinery, New York, NY, USA, Article 1, 10 pages. doi:10.1145/2537052.2537053
- [39] Cornershop Creative. 2018. UN Broadband Commission Adopts A4AI "1 for 2" Affordability Target. <https://a4ai.org/news/un-broadband-commission-adopts-a4ai-1-for-2-affordability-target/>
- [40] Leslie Daigle. 2004. *WHOIS Protocol Specification*. Request for Comments RFC 3912. Internet Engineering Task Force. doi:10.17487/RFC3912 Num Pages: 4.
- [41] Rodéric Fanou, Gareth Tyson, Pierre Francois, and Arjuna Sathiaselalan. 2016. Pushing the Frontier: Exploring the African Web Ecosystem. In *Proceedings of the 25th International Conference on World Wide Web* (Montréal, Québec, Canada) (WWW '16). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 435–445. doi:10.1145/2872427.2882997
- [42] Adrienne Porter Felt, Richard Barnes, April King, Chris Palmer, Chris Bentzel, and Parisa Tabriz. 2017. Measuring HTTPS Adoption on the Web. In *26th USENIX Security Symposium, USENIX Security 2017*, Vancouver, BC, Canada, August 16-18, 2017, Engin Kirda and Thomas Ristenpart (Eds.). USENIX Association, 1323–1338.
- [43] fido128. 2024. Answer to "CrUX: How much data is sufficient?". <https://stackoverflow.com/a/77872898>
- [44] Michelle WL Fong. 2009. Technology leapfrogging for developing countries. In *Encyclopedia of Information Science and Technology, Second Edition*. IGI Global, 3707–3713.
- [45] Google. 2019. AMP is a web component framework to easily create user-first web experiences - amp.dev. <https://amp.dev>. Accessed: 2019-05-05.
- [46] Matthias Gotze, Srđjan Matic, Costas Jordanou, Georgios Smaragdakis, and Nikolaos Laoutaris. 2022. Measuring web cookies in governmental websites. In *Proceedings of the 14th ACM Web Science Conference 2022*. 44–54.
- [47] World Bank Group. 2016. *World development report 2016: Digital dividends*. World Bank Publications.
- [48] GSMA. 2020. *The State of Mobile Internet Connectivity Report 2020*. Technical Report. GSMA. <https://www.gsma.com/r/somic/>
- [49] Rumaisa Habib, Aimen Inam, Ayesha Ali, Ihsan Ayyub Qazi, and Zafar Ayyub Qazi. 2023. A First Look at Public Service Websites from the Affordability Lens. In *Proceedings of the ACM Web Conference 2023*. ACM, Austin TX USA, 2731–2741. doi:10.1145/3543507.3583415
- [50] Rumaisa Habib, Sarah Tanveer, Aimen Inam, Haseeb Ahmed, Ayesha Ali, Zartash Afzal Uzmi, Zafar Ayyub Qazi, and Ihsan Ayyub Qazi. 2023. A Framework for Improving Web Affordability and Inclusiveness. In *Proceedings of the ACM SIGCOMM 2023 Conference* (New York, NY, USA) (ACM SIGCOMM '23). Association for Computing Machinery, New York, NY, USA, 592–607. doi:10.1145/3603269.3604872
- [51] Shaddi Hasan, Mary Claire Barela, Matthew Johnson, Eric Brewer, and Kurtis Heimerl. 2019. Scaling community cellular networks with community cellular manager. In *Proceedings of the 16th USENIX Conference on Networked Systems Design and Implementation* (Boston, MA, USA) (NSDI'19). USENIX Association, USA, 735–750.
- [52] Shaddi Hasan, Yahel Ben-David, Colin Scott, Eric Brewer, and Scott Shenker. 2013. Enhancing rural connectivity with software defined networks. In *Proceedings of the 3rd ACM Symposium on Computing for Development* (Bangalore, India) (ACM DEV '13). Association for Computing Machinery, New York, NY, USA, Article 49, 2 pages. doi:10.1145/2442882.2442937
- [53] Shaddi Hasan, Kurtis Heimerl, Kate Harrison, Kashif Ali, Sean Roberts, Anant Sahai, and Eric Brewer. 2014. GSM whitespaces: An opportunity for rural cellular service. In *2014 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 271–282. doi:10.1109/DySPAN.2014.6817804
- [54] Ellen Helsper. 2021. The digital disconnect: The social causes and consequences of digital inequalities. (2021).

- [55] Martin Hilbert. 2016. The bad news is that the digital access divide is here to stay: Domestically installed bandwidths among 172 countries for 1986–2014. *Telecommunications Policy* 40, 6 (2016), 567–581.
- [56] W3C Web Accessibility Initiative (WAI). [n. d.]. What's New in WCAG 2.2. <https://www.w3.org/WAI/standards-guidelines/wcag/new-in-22/>
- [57] Jordan Jueckstock, Shaown Sarker, Peter Snyder, Aidan Beggs, Panagiotis Papadopoulos, Matteo Varvello, Benjamin Livshits, and Alexandros Kapravelos. 2021. Towards realistic and reproducible web crawl measurements. In *Proceedings of the Web Conference 2021*. 80–91.
- [58] Wallace Koehler. 2002. Web page change and persistence - A four-year longitudinal study. *J. Assoc. Inf. Sci. Technol.* 53, 2 (2002), 162–171.
- [59] Zahir Koradia, Goutham Mannava, Aravindh Raman, Gaurav Aggarwal, Vinay Ribeiro, Aaditeshwar Seth, Sebastian Ardon, Anirban Mahanti, and Sipat Triukose. 2013. First impressions on the state of cellular data connectivity in India. In *Proceedings of the 4th Annual Symposium on Computing for Development* (Cape Town, South Africa) (*ACM DEV-4 '13*). Association for Computing Machinery, New York, NY, USA, Article 3, 10 pages. doi:10.1145/2537052.2537064
- [60] Raula Gaikovina Kula, Daniel M German, Ali Ouni, Takashi Ishio, and Katsuro Inoue. 2018. Do developers update their library dependencies? An empirical study on the impact of security advisories on library migration. *Empirical Software Engineering* 23 (2018), 384–417.
- [61] Jesutofunmi Kupoluyi, Moumena Chaqfeh, Matteo Varvello, Russell Coke, Waleed Hashmi, Lakshmi Subramanian, and Yasir Zaki. 2022. Muzeel: assessing the impact of JavaScript dead code elimination on mobile web performance. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) (*IMC '22*). Association for Computing Machinery, New York, NY, USA, 335–348. doi:10.1145/3517745.3561427
- [62] Tobias Lauinger, Abdelberri Chaabane, Sajjad Arshad, William Robertson, Christo Wilson, and Engin Kirda. 2017. Thou Shalt Not Depend on Me: Analysing the Use of Outdated JavaScript Libraries on the Web. In *24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February 26 - March 1, 2017*. The Internet Society.
- [63] Kyunchan Lim, Yonghwi Kwon, and Doowon Kim. 2023. A Longitudinal Study of Vulnerable Client-side Resources and Web Developers' Updating Behaviors. In *Proceedings of the 2023 ACM on Internet Measurement Conference, IMC 2023, Montreal, QC, Canada, October 24-26, 2023*, Marie-José Montpetit, Aris Leivadeas, Steve Uhlig, and Mobin Javed (Eds.). ACM, 162–180.
- [64] Sophie Lythreathis, Sanjay Kumar Singh, and Abdul-Nasser El-Kassar. 2022. The digital divide: A review and future research agenda. *Technological Forecasting and Social Change* 175 (2022), 121359.
- [65] Usama Naseer and Theophilus A. Benson. 2023. JS Capsules: A Framework for Capturing Fine-grained JavaScript Memory Measurements for the Mobile Web. *Proc. ACM Meas. Anal. Comput. Syst.* 7, 1, Article 14 (March 2023), 26 pages. doi:10.1145/3579327
- [66] Usama Naseer, Theophilus A. Benson, and Ravi Netravali. 2021. WebMedic: Disentangling the Memory-Functionality Tension for the Next Billion Mobile Web Users. In *Proceedings of the 22nd International Workshop on Mobile Computing Systems and Applications* (Virtual, United Kingdom) (*HotMobile '21*). Association for Computing Machinery, New York, NY, USA, 71–77. doi:10.1145/3446382.3448652
- [67] OECD. 2018. *Bridging the Digital Gender Divide*. Technical Report. Organisation for Economic Co-operation and Development. <https://www.oecd.org/digital/bridging-the-digital-gender-divide.pdf>
- [68] Ihsan Ayyub Qazi, Zafar Ayyub Qazi, Ayesha Ali, Muhammad Abdullah, and Rumaisa Habib. 2021. Rethinking Web for Affordability and Inclusion. In *Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks*. ACM, Virtual Event United Kingdom, 9–15. doi:10.1145/3484266.3487376
- [69] Jeremy Rack and Cristian-Alexandru Staicu. 2023. Jack-in-the-box: An Empirical Study of JavaScript Bundling on the Web and its Security Implications. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 3198–3212.
- [70] Jannis Rautenstrauch, Metodi Mitkov, Thomas Helbrecht, Lorenz Hetterich, and Ben Stock. 2024. To Auth or Not To Auth? A Comparative Analysis of the Pre- and Post-Login Security Landscape. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1500–1516.
- [71] Ali Raza, Yasir Zaki, Thomas Pötsch, Jay Chen, and Lakshmi Subramanian. 2017. xCache: Rethinking Edge Caching for Developing Regions. In *Proceedings of the Ninth International Conference on Information and Communication Technologies and Development* (Lahore, Pakistan) (*ICTD '17*). Association for Computing Machinery, New York, NY, USA, Article 5, 11 pages. doi:10.1145/3136560.3136577
- [72] E. Rescorla. 2000. HTTP Over TLS.
- [73] Sebastian Roth, Lea Gröber, Michael Backes, Katharina Krombholz, and Ben Stock. 2021. 12 Angry Developers - A Qualitative Study on Developers' Struggles with CSP. In *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*. ACM, 3085–3103.
- [74] Kimberly Ruth, Aurore Fass, Jonathan Azose, Mark Pearson, Emma Thomas, Caitlin Sadowski, and Zakir Durumeric. 2022. A world wide view of browsing the world wide web. In *Proceedings of the 22nd ACM Internet Measurement Conference, IMC 2022, Nice, France, October 25-27, 2022*, Chadi Barakat, Cristel Pelsser, Theophilus A. Benson, and David R. Choffnes (Eds.). ACM, 317–336. doi:10.1145/3517745.3561418
- [75] Asheesh Sharma, Manveen Kaur, Zahir Koradia, Rahul Nishant, Sameer Pandit, Aravindh Raman, and Aaditeshwar Seth. 2015. Revisiting the State of Cellular Data Connectivity in India. In *Proceedings of the 2015 Annual Symposium on Computing for Development* (London, United Kingdom) (*DEV '15*). Association for Computing Machinery, New York, NY, USA, 149–157. doi:10.1145/2830629.2830649
- [76] Peter Snyder, Cynthia Taylor, and Chris Kanich. 2017. Most websites don't need to vibrate: A cost-benefit approach to improving browser security. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 179–194.
- [77] International Telecommunication Union. 2020. *Measuring digital development: ICT Price Trends 2020*. Technical Report. ITU. <https://www.itu.int/en/ITU-D/Statistics/Pages/default.aspx>
- [78] Yasir Zaki, Jay Chen, Thomas Pötsch, Talal Ahmad, and Lakshminarayanan Subramanian. 2014. Dissecting Web Latency in Ghana. In *Proceedings of the 2014 Conference on Internet Measurement Conference* (Vancouver, BC, Canada) (*IMC '14*). Association for Computing Machinery, New York, NY, USA, 241–248. doi:10.1145/2663716.2663748
- [79] Yasir Zaki, Jay Chen, Thomas Pötsch, Talal Ahmad, and Lakshminarayanan Subramanian. 2014. Dissecting Web Latency in Ghana. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, Vancouver BC Canada, 241–248. doi:10.1145/2663716.2663748
- [80] Eugenia Zigisova. 2022. *The 2022 Web Almanac: Third Parties*. Technical Report. HTTP Archive. <https://almanac.httparchive.org/en/2022/third-parties> Issue: 8 Publication Title: The 2022 Web Almanac Volume: 4.

A Appendix

A.1 Distribution of Website Rankings

Figure 16 displays the distribution of website rankings across different countries in our dataset. The figure illustrates that, for most countries, the 10,000 sampled websites fall within the top 50,000 of that country, indicating a focus on relatively popular websites. However, for countries such as Bangladesh, Pakistan, Nigeria, and the Philippines, less popular websites were included to meet the target sample size of 10,000 websites per country.




Figure 16: Distribution of website rankings in our dataset per country. The figure shows that in most cases, the 10,000 websites per country are within the top 50,000 of that country. For Bangladesh, Pakistan, Nigeria, and the Philippines, less popular websites were included to meet the target sample size.

A.2 Country-wise Distribution of Image Sizes

Figure 17 illustrates the distribution of image sizes across different countries. The chart highlights that websites in developing countries tend to use larger image sizes compared to those in developed countries. Specifically, 15% of images on websites in developing countries exceed 100 KB, whereas only 12% of images in developed countries fall into this category. This data supports the observation that while websites in developing countries may not host more

images overall, they often feature larger images, contributing to greater webpage sizes and potentially increased load times and bandwidth usage.




Figure 17: Distribution of image sizes across countries.

A.3 Distribution of Document Sizes

Figure 18 illustrates the distribution of document sizes across different countries. The data indicates that websites in developed countries tend to have larger document sizes compared to those in developing countries.




Figure 18: Distribution of document sizes across countries.

A.4 Country-wise Distribution of Total Blocking Time

Figure 19 illustrates the distribution of total blocking time across developed and developing countries. Total blocking time refers to the duration during which the browser’s main thread is blocked, preventing user interactions such as clicks, scrolling, or typing. This metric is important for assessing website responsiveness, with longer blocking times leading to poorer user experiences. The data shows significant differences between regions, with developed countries generally exhibiting longer blocking times compared to developing countries. This disparity is largely attributed to the heavier use of JavaScript in websites from developed regions, which increases the complexity of web pages and leads to more frequent blocking of the main thread.




Figure 19: Distribution of total blocking time across countries.

A.5 HTTPS Adoption




Figure 20: Country-wise HTTPS and HTTP adoption rates.

Figure 20 shows significant variation in HTTPS adoption across countries. Developed nations like the United States and Germany have higher adoption rates, while developing countries like Bangladesh and China exhibit a larger percentage of websites still relying on HTTP. This suggests that security practices vary significantly between regions, potentially exposing users in some areas to greater security risks.

A.6 CSS Wastage Distribution




Figure 21: Percentage of unused CSS rules across countries

Figure 21 presents the percentage of unused CSS rules across various countries. Developing countries such as Pakistan and Bangladesh exhibit higher wastage levels, exceeding 95%, while developed countries like Japan and Germany show slightly lower wastage rates, around 92%.

A.7 PAW Analysis

The PAW (Price Adjusted Web access) Index is a measure that evaluates the affordability of accessing web content in a given region by comparing the average broadband prices, local webpage sizes, and income levels to international affordability targets. The formula for the PAW Index is:

$$PAW_i = \frac{P_i}{P_T} \times \frac{W_{j,i}}{W_{global}}$$




Figure 22: Total accessibility violations across countries. The similar median values suggest that accessibility challenges are a global issue.

where P_i represents the average broadband price in country i , P_T is the target broadband price set at 2% of the country’s GNI per capita, $W_{j,i}$ is the size of webpage j in country i , and W_{global} is the global average webpage size. A PAW Index greater than 1 indicates that accessing web content is unaffordable in comparison to the international benchmark, while a value less than or equal to 1 suggests that the region meets the affordability target. This index provides a tool for policymakers and researchers to assess and compare the digital affordability landscape across regions, helping to identify gaps and make informed decisions to promote accessible digital content.

B Accessibility Violations Analysis

Figure 22 shows the total number of accessibility violations across countries, highlighting that accessibility challenges are a global issue affecting both developed and developing nations. Similar trends have been observed in prior studies [30], which indicate that accessibility issues persist worldwide due to a lack of prioritization and awareness, rather than economic or technological limitations. To assess each website’s accessibility, we use Google Lighthouse, an automated tool that evaluates web content against the Web Content Accessibility Guidelines (WCAG) [21]. Lighthouse calculates

an Accessibility score as a weighted average of audits, informed by axe-core [1], an open-source engine that analyzes user impact and identifies compliance issues. Our evaluation specifically targets WCAG 2.2 [56], which is organized around four principles: Perceivable, Operable, Understandable, and Robust (POUR). These principles ensure accessibility for diverse user needs, including those with disabilities, across various devices.

The median number of accessibility violations remains consistent at 4 for most countries, with exceptions like Bangladesh, Korea, and Russia recording slightly higher medians of 5. However, this uniformity belies significant disparities in maximum violation counts: Bangladesh leads with 14 violations, followed by Korea (13), and Australia, the Philippines, and the United Kingdom (12). Most nations, including Brazil, China, Germany, France, India, and Pakistan, cluster around 10 violations, while Canada, Indonesia, Nigeria, and the United States report marginally lower counts at 9. The striking similarity in median values across diverse economic and technological contexts underscores that accessibility challenges are systemic rather than isolated to specific conditions. This pattern suggests violations stem not from national wealth or infrastructure gaps but from a pervasive lack of global awareness and prioritization of accessibility standards in development practices.

These findings emphasize the urgent need for coordinated action to foster inclusivity. Prioritizing developer education, enforcing compliance incentives, and harmonizing accessibility policies globally could mitigate barriers and create a more equitable digital landscape for all users.

C Dataset Availability

The dataset used in this study, including a summary and detailed descriptions, is available on GitHub [4]. Due to its large size, the full dataset is not hosted directly on GitHub but can be provided upon request. For more details on the dataset structure, access instructions, and request procedures, please refer to the GitHub repository.