Towards Browser-Based Consent Management

Gayatri Priyadarsini Kancherla
Indian Institute of Technology, Gandhinagar, India
gayatripriyadarsini@iitgn.ac.in

Abhishek Bichhawat
Indian Institute of Technology, Gandhinagar, India
abhishek.b@iitgn.ac.in

***Abstract*—Today's users are concerned about the privacy of their personal or sensitive information on the Web because of the different techniques being employed to track their activities and behaviour online. Privacy laws like the GDPR, CCPA, etc., provide some control to the user to decide whether they would like to share their personal data online and what and how much they are willing to share. These laws require the websites to be transparent to the users about what information they are collecting and how that information shall be used, and insist that the websites obtain explicit consent from the users before collecting this information. However, the consent given by the users may not always be honoured by the websites, by tracking the user despite their rejection of advertisement and analytics cookies. Additionally, recent studies in the area also show that websites often utilize *dark patterns* through manipulative design choices. This affects the consent choices of the users, thereby tricking them into consenting to share more information than what they actually intend.**

We propose an alternate consent management system that shifts the trust from the web servers to browsers, i.e., instead of relying on servers to obtain and comply with the consent provided by the user, we delegate this task to the web browser. In our approach, the browser obtains and stores the consent of the user for the visited websites through a standard consent banner in the user's preferred language, irrespective of the website's language. The cookies set by the websites are then subject to this consent provided by the user, as each of the cookies carries an additional attribute that identifies their category. This approach provides an easier way for users to manage consent for different websites without having to search for policies and compliance with the websites while also solving the language barrier. We modified the Nightly Firefox build to integrate an additional cookie attribute that stores the purpose of the cookie, a consent banner to get the user's cookie preferences and implement the required checks during cookie access.

We believe this approach offers a streamlined and more transparent methodology for managing consent, making it easier to audit and maintain.

I. Introduction

In recent years, users have become more aware of the privacy of the data that they share on the internet. Many privacy laws like GDPR [24], CCPA [15], etc., have come into effect in different regions of the world to ensure that the tracking and sharing of user data happen with the permission and knowledge of the users. The websites need to obtain explicit consent from the users before they collect personal data or store tracking information other than what is required for the proper functioning of the website. They are also obliged to share with users the different ways in which their data will be used. Websites normally use consent banners to obtain user permissions and store tracking data in cookies, localstorage, etc. In the current framework, web servers are trusted to respect the consent provided by the users.

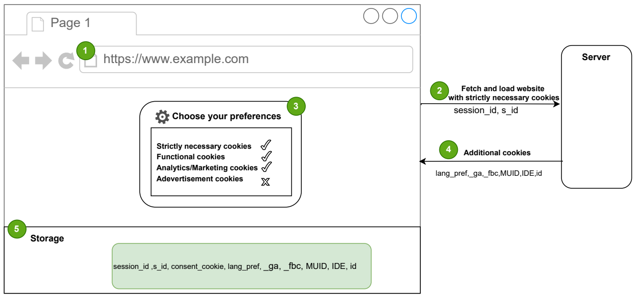
Page 1
https://www.example.com
1
2
Fetch and load website with strictly necessary cookies
session_id, s_id
3
Choose your preferences
Strictly necessary cookies ✓
Functional cookies ✓
Analytics/Marketing cookies ✓
Adevertisement cookies ✗
4
Additional cookies
lang_pref,_ga,_fbc,MUID,IDE,id
Server
5
Storage
session_id ,s_id, consent_cookie, lang_pref, _ga, _fbc, MUID, IDE, id

Fig. 1: Current Consent Management Framework

Figure 1 illustrates the current framework used by websites. Suppose a user visits a website, example.com, from their browser ①. The browser sends a request to the server to retrieve the web page, and the server responds to the request along with the strictly necessary cookies required for the proper functioning of the website ②. The webpage also shows a consent banner to the user to obtain their permission ③. Suppose the user selects functional, analytics, and marketing cookies but not advertisement cookies in addition to the strictly necessary cookies. These choices, when shared with the server, result in additional cookies belonging to the selected categories being sent by the server ④. The major concern with this approach is that it is not clear whether the servers and the third-party scripts on a web page abide by their privacy policy and consent provided by the user. In other words, it is possible that the server may set more cookies (at ⑤) than were agreed upon by the user, and the user does not have an easy way of checking if their consent is respected.

The absence of clear regulations and standardized implementation practices complicates the auditing process, making it difficult to identify potential violations and ensure compliance. The current consent management framework faces several challenges and limitations:

1) Previous studies [26], [43] have shown that the text and design used for informing the user is not consistent and may not abide by the legal regulations. Ideally, the list of cookies stored in the browser should match those listed

Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) 2025
28 February 2025, San Diego, CA, USA
ISBN 979-8-9894372-9-0
https://dx.doi.org/10.14722/madweb.2025.23021
www.ndss-symposium.org

in the website's privacy policy, but this information is not always easily accessible.

2) Websites may not be available in the user's preferred language, and hence the consent banner's text may not be comprehensible to the user.
3) Prior works have identified dark patterns used by websites or consent management platforms (CMPs) to trick users into consenting to share more information than intended [9], [10], [12], [25], [27], [36], [41], [45].
4) Since there is no standard way to identify whether the cookies belong to a specific category, it is difficult to identify non-compliance when optional cookies are set even after consent is not provided. Recent studies have shown that websites often set more cookies than users consent to, and not all cookies stored in the browser match those listed in the privacy policy [7], [11], [13], [16], [28], [44]. However, these methods are not completely accurate in identifying the category of the cookies, leading to either functionality breakage of the website or leaving them vulnerable to bypassing [20].
5) The representation of the user's consent and storing the same in the browser is not standardized. The websites are free to use any format to use and share the user's consent. While frameworks like [8] and [3] have been suggested and recommended by regulators, not all websites use these frameworks.

The goal of our work is to design an alternate unified framework for consent management that provides better transparency to the user and easier adaptability and legal compliance for the websites.

Recent work by Bruhner et al. [14], which built on the Advanced Data Protection Control [29], proposed moving the responsibility to the browser by providing guidelines for how such an enforcement on the browser side may work. The idea is to delegate the task of consent management to the browsers that, based on the preferences of the users, allow or deny cookies to be set by websites. We further build on this work providing concrete action items for putting such a system in place. In this work, we develop an approach to standardize user consent management such that the *browsers* ensure consent compliance for a website.

The advantages of this approach are manifold: (1) browsers can ensure that the users' inputs are processed correctly without relying on an extension (2) this removes the language barrier, which makes it difficult for certain users to comprehend a consent banner across boundaries (3) there is *accountability* for the websites when tagging the cookies with specific purposes, which is not available in the current setting.

To enforce this, we propose an additional cookie attribute that records the purpose of each of the cookies. This would require the server to specify the category of the cookie as a cookie field while creating it. This would be stored and used at the browser end to decide if a particular cookie is allowed in a session.

Additionally, the browser exposes an option that allows the user to select different categories of cookies that they consent to. Based on the choices selected by the user, the browser enforces checks on the cookies and removes those whose purpose does not match the choices. To ensure that the third-party scripts do not misuse the attribute, we enforce that only the host can set cookies that are tagged as strictly necessary. Similarly, we allow only the first-party or host scripts to modify this attribute of the cookie.

As a prototype, we instrument the Nightly Firefox build to process the additional cookie attribute when set by the host page or the scripts and store it in a *protected* manner in the cookie jar of the browser. In the default setting, we are currently implementing the interface to expose an option to the user in the browser for selecting their choices. We store the user's consent as a browser preference and use the value to compare it with the category of the cookie being set.

Our framework provides a streamlined way to manage consent for the users while providing trust and accountability. Next, we discuss some of the related works in the area that inspired our approach, followed by the detailed methodology and the planned future work.

II. Related Works

A. Cookie classification

In this section, we discuss the works that have evaluated the current scenario of consent management. Previous works that have tried to classify cookies into categories have had to depend on either privacy policies, cookie tables, or machine learning and natural language processing techniques. However, these efforts either do not cover every website's declaration of cookies or are not entirely accurate in classifying the cookies. Some of the popular approaches of handling the cookie banners at the client side include Consent-O-Matic [17], [36] and Autoconsent [35] that depend on the presence of particular CMPs to handle consent at the user end. However, they can only be used on websites that use the supported CMPs.

Hu et al. [28] propose a cookie purpose classifier using Naive Bayes classifier; they used Cookiepedia [38] as their training data and could only achieve an F1 score of 83%. Similarly, in the work by Calzavara et al. [16], their model's accuracy was 83% for cookie classification. Bollinger et al. [11] give a client/user side solution to enforce GDPR by deleting the cookies that are not consented to by the user. They use machine learning to classify the cookies, which are then deleted based on the consent provided by the users. However, the classification may not be extensible to all websites since their training data only contained data from policies of websites from three CMPs. While we propose a similar client-side solution, our approach requires explicit labeling by the server, thereby avoiding the (machine) learning phase altogether.

B. User studies and dark patterns

Recent user studies have shown that users find it difficult to manage privacy on their end because of the long and incomprehensible privacy policies and cookie banners. The study by Alharbi et al. [9] evaluates user perspectives on privacy and security by assessing cookie interfaces of e-government websites from 50 countries across Europe, America, Oceania, and Asia, using individual expert review methods. Their research highlights significant usability issues, as many websites fail to meet privacy guidelines, resulting in poor user awareness and management of privacy settings. Over

90% of the websites were found to use dark patterns, which mislead users and complicate their privacy choices. The study concludes with recommendations for designing user-friendly and GDPR-compliant cookie interfaces to enhance user trust and control over personal data.

Various prior works [10], [12], [25], [27], [36], [41], [45] discuss the existence of *dark patterns* [23] adopted by websites in order to manipulate the users into accepting the cookies in their browser. Despite the presence of well-defined regulations like GDPR, Nouwens et al. [36] showed that 88.2% out of 680 examined websites that use CMPs violate the simple requirements mentioned by GDPR. Similarly, other works by Bielova et al. [10] and Gray et al. [25] also study the effects of dark patterns and their prevalence on websites, with a focus on the legal specifications as well. Some other studies that focus on dark patterns include [12], [27], [41], [45].

Additionally, the security implications discussed in Klein et al. in their work [34] regarding cookie banners indicate the implications of accepting consent cookie banners. They show that accepting consent increases the third-party scripts by 45% and exposes them to a 63% increase in information flow on average. They further show that their XSS exploits worked on 55% websites when consent is given, indicating the serious security vulnerabilities added to a website when third-party scripts are added due to default acceptance being given to a website. Understanding the security implications of cookie banners further reinforces the importance of designing cookie consent notices that facilitate informed decision-making. Bouhoula et al. [13] present a tool that uses NLP to identify non-compliance, which reports that 65.5% of the websites offering a rejection option likely collect user data despite explicit negative consent.

C. Legal aspects of enforcement

Another side to this is that the legal aspects of the privacy laws and acts are not very specific and direct for website developers to follow. In their work, Santos et al. [42] studied the effectiveness of GDPR, where they provide specific requirements for developers and regulators to ensure privacy and discuss verification challenges. They also address inconsistencies in policies among different Data Protection Authorities (DPAs) and the need for standardization. Similarly, Degeling et al. [19] conducted a measurement study and observed a 16% increase in the display of cookie consent interfaces among 6,579 evaluated websites after the GDPR came into effect. The Interactive Advertising Bureau Europe's (IAB Europe) Transparency and Consent Framework (TCF) was subsequently developed as a standardized approach by industry stakeholders to manage user consent and preferences for data processing and targeted advertising. Complications arising from its deployment led the developers to depend on different consent management platforms (CMPs) to abide by the regulations. Toth et al. [44] emphasize that website publishers need to monitor user-data processing and consent management and not leave it to the CMPs. They discuss how the CMPs may not comply with the law and use deceptive design schemes.

D. Alternate frameworks for consent management

The Platform for Privacy Preference Project (P3P) [46] is a framework proposed to ease the visualization of privacy policies at the user-end for taking user consent and enforcing the user's preferences at the server end. The Privacy Bird [18] was one of the implementations of this framework, and extensions include proposals like [37], [40]. However, these approaches were not adopted. These frameworks involved scanning through privacy policies, mapping user preferences to the HTML input elements, etc. These also involved cooperation between multiple parties and difficult-to-implement changes at both the client and server-side due to the unambiguous nature of the privacy policies or banners and consent banners provided by the websites.

Global Privacy Control [2] is a browser-based solution to signal the websites to indicate that the user does not want to be tracked. Zimmeck et al. [47] in their study discuss how DNT signal [1] and other consent storing standards like IAB provided USP flag [3] can be used to check its enforcement. Few studies [5], [39] discuss how the Do Not Track (DNT) flag in the header is not respected by the websites, which further highlights the need for a browser-based solution. Our solution will not leave it on the website to respect the users' choice and will do that on the browser end.

Another such technical specification proposed in recent years is the Advanced Data Protection Control (ADPC) [29]. The motivation was to have bidirectional communication between the user and the server by sending the user's consent to the server as part of the header. Additionally, interaction with the banners becomes easier by interacting with the extension that abides by the legal regulations. However, consent and cookie management are still being done on the server side in this case. The server is *expected to honor the user's consent* and set the cookies the user has consented to.

Bruhner et al. [14] propose a framework extending the ADPC as ADPC+, which moves this responsibility to the browser. The framework suggests five requirements — (1) no cookies are stored before the user's consent is given, (2) all the information to be given by the data controller will be available at a specific location for the browser to fetch and enforce on the browser side, (3) user will be able to change their consent, (4) update in the user's consent will be communicated to the server through the header, and (5) cookies other than ones classified as necessary, require user's consent. We further build on these suggestions, providing concrete enforcement strategies to realize the framework's goal. Instead of fetching the privacy information from the server at the run-time, we propose to get the purpose of each cookie as an attribute itself. Additionally, such labeling can be extended to other storage objects as well, which have been found to store consent-related information as well in prior studies [32], [33]. We believe that our development would further the potential applicability of the abstract model from Bruhner et al. [14], showcasing the benefits of this approach.

III. Methodology

In this section, we describe our proposed framework. Recall that the server plays an important role in managing consent in the current framework (Figure 1). It is responsible for obtaining consent from the user (sometimes via a CMP), recording their preferences, and, accordingly, setting the cookies on the client side based on the mapping defined in their cookie/privacy

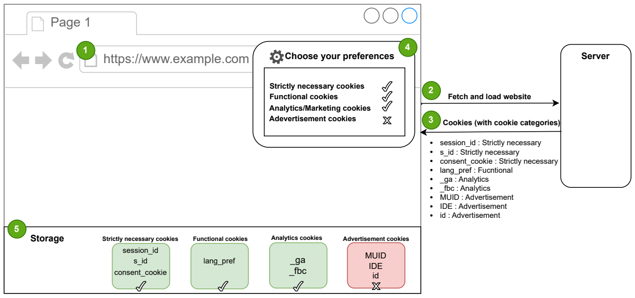

Fig. 2: Proposed consent management framework

policy. However, as prior works have shown, servers may not be handling the users' consent as per their policy [11], [13], [25], [44].

Our proposed framework attempts to standardize the approach by handling consent on the client side (browser). Figure 2 illustrates the different steps involved when handling consent on the browser side.

(1) User visits the website.
(2) Browser sends a request to the server.
(3) Server sends the response along with all the cookies **labeled with categories** as an additional attribute.
(4) Browser asks the user for their preference.
(5) Based on the user preference, selected cookies are stored at the client. The preference is also recorded for future visits.

While our approach shifts the responsibility of handling consent to the browser, it still relies on the server to decide the category associated with the cookies, thereby ensuring that the functionality of the website is not broken. This is in line with the current architecture where the server decides which cookies to keep in the browser based on the user preferences and the mapping in the privacy policy.

Next, we discuss the main aspects of our architecture and the proposed modifications.

A. *Server-side modifications*

The server's main task would be to add an additional cookie field/attribute when setting the cookies on the website. This makes the adoption easier since there is no need to deploy a CMP explicitly to register and manage the user consent. Similarly, third-party scripts on the page, which query the CMPs for consent in the current scenario, can now query APIs provided by the browser to access the users' consent.

For example, a strictly necessary cookie with the additional field 'Category' would be: "sess_id=12345; path=/; domain=hostdomain.com; secure; HttpOnly; SameSite=Strict; **Category=1000**".

B. *Browser-side changes*

We modified the Nightly Firefox build to partially enforce the suggested browser modifications. We describe the changes done so far and their extensions.

a) New cookie attribute: To add an additional cookie field or attributes and the required checks for using this attribute, we modified the files related to handling cookies. In Firefox, we add 150 lines across six primary files that handle cookie-related services. We store the category along with every cookie, which is a four-digit integer. Each digit represents one of the four categories, 'necessary', 'functional', 'analytics', and 'advertisement', in this order, that the cookie can belong to.

We selected these four categories following the established classifications recommended by GDPR [6] and UK ICC [4] cookie guidelines. While these categories were sufficient for our proof-of-concept implementation, our framework is designed to be extensible. The bit-based category representation can be expanded up to 64 bits to accommodate additional categories as defined in more comprehensive standards like IAB Europe's Transparency and Consent Framework (TCF) [30].

By default, the value of the category attribute of the cookies is set to 0000. The default value would block access to the cookie since it does not belong to any category. For example, category=0010 indicates that the cookie is used for analytics by the website. This format of the attribute will make it easier for the cookie to belong to multiple categories. For instance, if a functional cookie is used for analytics for the website's proper functionality, the website may set the attribute as category=0110. A cookie marked with multiple categories will only be accepted if both those categories are accepted by the user. However, necessary cookies are allowed by default, irrespective of the other categories. Access to cookies marked with no category (i.e., 0000), by default, will be blocked.

We added this additional cookie field to the cookie structure and modified CookiePersistentStorage.cpp, which handles the database storage operations. The responses are handled in the CookieService.cpp and CookieServiceChild.cpp, which ensure that only user-consented cookies are sent in the response. The checks made in these files include (1) change to the category field of a cookie, which will only be allowed by a script *belonging to the host* domain. Similarly, we want to (2) refrain the third-party scripts from being able to set strictly necessary cookies and limit it to the host page scripts.

b) User interface: We create a built-in extension that injects a cookie banner on the webpage. This banner provides the following features:

- Simple and concise design and language.
- Check boxes for each category for fine-grained selection.
- Explanation of each category.
- Accept all and Reject all buttons
- Drop down to change language preference.

c) User consent preference: We save the user's consent preferences as part of the browser preferences. Consent is stored *for every website* in the form of 4 integers, similar to the proposed cookie attribute. For example, consent_object=1010 indicates that the user has consented to necessary and analytics cookies. This object will be used to make checks when the cookies are accessed. By default, the value of this object is set to allow necessary cookies, i.e., 1000, and the category attribute of the cookies is set to 0000. The default value would block access to the cookie since it does not belong to any category.

The user can provide their language preference through the privacy and security preferences page. This would then be useful to later the cookie banner in that language.

This page also includes the list of domains and the corresponding consent given by the user. This can further be visited by the user to change for future visits.

d) Consent update event listener: If a user accepts advertisement and analytics cookies, third-parties receive these cookies. Regulations [15], [21], [22], [24], [31] require these third-parties to stop processing on the user data if the user decides to later change their consent. This is usually achieved by the CMPs by providing an event listener on the button clicks on the consent banner.

Since the user can change their preferences at any point, we additionally propose to add an API for accessing the user consent storing string to obtain the consent object's value whenever it changes. This way, the scripts can be notified of the updated consent.

IV. Preliminary Evaluation

Our implementation establishes user consent preferences as a browser-level configuration maintained within a dedicated consent_object. This approach enables systematic consent validation for cookie operations. When any cookie access is attempted, the system automatically compares the cookie's declared category against the stored user preferences in the consent_object.

To validate our implementation, we developed a comprehensive test environment consisting of a demo website specifically designed to create and manipulate cookies with explicit category assignments. We conducted systematic testing using our instrumented browser implementation , focusing on the following critical security and functionality aspects:

1) Access to cookies should be blocked when the category is not allowed by the user.
2) Access to cookies should be blocked when the category attribute is not set by the server while creating cookies.
3) Script domain that is not the creator should not be able to write the category attribute of the cookie.

Our framework successfully addressed all test scenarios, demonstrating effective browser-side consent management. The testing process also revealed several implementation constraints and potential areas for enhancement, which we examine in detail in Section VI. These findings include considerations for scaling the solution across diverse web applications, handling legacy cookies, and potential performance optimizations.

V. Limitations

a) Malicious Cookie Category Classification: A significant security concern arises when servers deliberately misclassify cookie categories, particularly by designating non-essential cookies as "necessary" to circumvent user consent mechanisms and ensure default browser acceptance. This vulnerability persists in the current scenario. However, our proposed framework introduces accountability by maintaining an immutable record of declared cookie purposes within the browser's storage mechanism.

b) Implementation Requirements for Novel Cookie Attributes: The incorporation of additional cookie attributes necessitates modifications to server-side implementations, requiring developers to update their existing codebase. This requirement aligns with historical precedents of security attribute implementations, such as the SameSite attribute, where similar modifications were necessary to enhance security measures. The established pattern of adopting new security attributes demonstrates the feasibility and acceptability of such implementation requirements.

c) Other consent storage mechanisms: While prior research has explored storing consent preferences in local storage [9], extending our cookie-based consent mechanism to local storage presents unique challenges. Unlike cookies, which possess a structured format with well-defined attributes, local storage implements a simple key-value pair model accessed through basic setter and getter functions. This architectural difference necessitates a different approach to consent management. We propose two potential solutions: (1) introducing new specialized setter and getter functions that incorporate consent categorization as part of the storage operation, or (2) modifying the existing API functions to mandatorily include consent information. This adaptation is crucial to maintain consistent consent management across different client-side storage mechanisms.

d) Essential third-party cookies: By default, we don't provide the option for third-party cookies to be marked as necessary for security reasons. However, SSO and similar authentication cookies are set by third-party but are often essential for functionality, and blocking them based on domain-level consent could break critical website features. Possible solutions for the same could be to (1) add a special category for essential third-party services or (2) white list SSO providers like Google, Facebook, and Microsoft at the browser level.

VI. Discussion

The main goal of our work is to build a robust framework for consent management, addressing certain shortcomings of the current model, which we discuss next.

a) Mismatch between cookies and policy: Previous studies have shown that consent management by servers or CMPs is often poorly enforced, with more cookies being set than those accepted by the user. To verify this mismatch, users must manually check the privacy policy for cookie purposes, which most don't do. Although some solutions intercept and modify cookies based on consent, they rely on data analysis or machine learning, leading to false positives and negatives. Websites can bypass these methods once they learn the heuristics used.

Our proposed approach tags the cookies with the purpose listed in the privacy policy of the webpage, providing an accurate means of managing cookies. As is already possible with the current model, the servers can still misinform the browser about the purpose of the cookie by incorrectly labeling the cookie's purpose. However, the server **can be held**

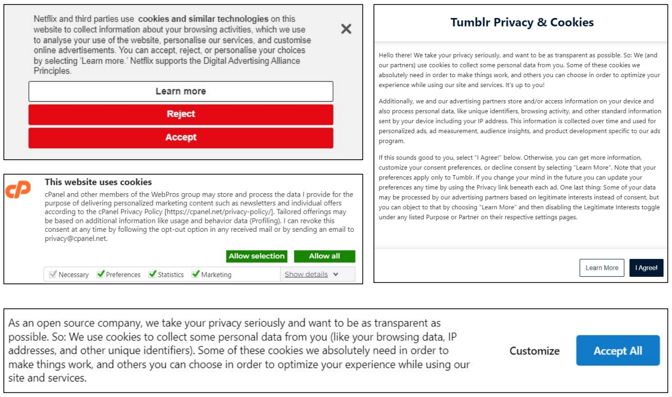

Netflix and third parties use **cookies and similar technologies** on this website to collect information about your browsing activities, which we use to analyse your use of the website, personalise our services, and customise online advertisements. You can accept, reject, or personalise your choices by selecting 'Learn more.' Netflix supports the Digital Advertising Alliance Principles.

Tumblr Privacy & Cookies

Hello there! We take your privacy seriously, and want to be as transparent as possible. So: We (and our partners) use cookies to collect some personal data from you. Some of these cookies we absolutely need in order to make things work, and others you can choose in order to optimize your experience while using our site and services. It's up to you!

Additionally, we and our advertising partners store and/or access information on your device and also process personal data, like unique identifiers, browsing activity, and other standard information sent by your device including your IP address. This information is collected over time and used for personalized ads, ad measurement, audience insights, and product development specific to our ads program.

If this sounds good to you, select "I Agree!" below. Otherwise, you can get more information, customize your consent preferences, or decline consent by selecting "Learn More". Note that your preferences apply only to Tumblr. If you change your mind in the future you can update your preferences any time by using the Privacy link beneath each ad. One last thing: Some of your data may be processed by our advertising partners based on legitimate interests instead of consent, but you can object to that by choosing "Learn More" and then disabling the Legitimate Interests toggle under any listed Purpose or Partner on their respective settings pages.

This website uses cookies

cPanel and other members of the WebPros group may store and process the data I provide for the purpose of delivering personalized marketing content such as newsletters and individual offers according to the cPanel Privacy Policy [https://cpanel.net/privacy-policy/]. Tailored offerings may be based on additional information like usage and behavior data (Profiling). I can revoke this consent at any time by following the opt-out option in any received mail or by sending an email to privacy@cpanel.net.

☑ Necessary ☑ Preferences ☑ Statistics ☑ Marketing

As an open source company, we take your privacy seriously and want to be as transparent as possible. So: We use cookies to collect some personal data from you (like your browsing data, IP addresses, and other unique identifiers). Some of these cookies we absolutely need in order to make things work, and others you can choose in order to optimize your experience while using our site and services.

Fig. 3: Consent banners on different websites

accountable for setting the purpose as the cookie is stored as part of the browser.

b) Standardized means for consent collection: Every website has a different means for collecting consent from the user (Figure 3). While some of them show their own consent banners to collect consent from the users, some other websites employ consent management platforms (CMPs) for managing consent. Moreover, there is no standard banner that a user interacts with across different websites, making it a difficult and unpleasant experience for the user. Additionally, these banners may be displayed inconsistently to users depending on their geographic locations. With the consent management delegated to the browser, the user would view a standard option for providing consent, thereby reducing the confusion and complications that would have arisen across multiple interfaces.

c) Consent in private modes: Another challenge with consent management is the lack of access to consent when the private or incognito mode in the browser is used. This makes it a repetitive task for the user every time they access a webpage, requiring them to provide their consent on every access. With the task delegated to the browser, the consent can be recorded once and used across all modes in the browser without having to obtain the input from the user.

d) Language barrier with consent banners: A major issue with the current framework is the language barrier when accessing websites in different countries. For instance, many of the consent banners for users in Germany are shown in German; users who may not have knowledge of the language require a translator to understand the text and the options. This makes it difficult for the user to check the right option without having the translation in place. The browser, on the other hand, can display the options in a language that the user understands, making the process of providing consent easier for the user.

e) Functionality Breakage: By default, the value of the category attribute for a cookie is set to 0000, blocking any access to this cookie. While this will ensure that the category attribute is used properly, this may result in functionality breakage if a legitimate cookie was mistakenly not categorized. For cases like this, we propose to alert the server in a separate header that contains these cookies and who tried to access them. Orthogonally, we are collecting the data pertaining to the use of strictly necessary cookies by third-party scripts and determining the breakage that it may cause in the current framework.

VII. Future Plans

We have currently implemented the parsing of additional cookie attributes on the browser side to show the possibility of enforcing the checks. We tested the implementation with a server sending the custom attributes and showed that the attribute is properly handled in the browser. As our next steps, we would implement an interface in the form of a button or an extension that allows the user to specify their consent options on a per-website basis and/or across all websites. We also need to enforce the rules on JavaScript access of consent and expose necessary interfaces for scripts to access the users' consent from the browser. Another interesting direction for future work is to explore the possibility of extending this to localstorage objects, which are used by some of the websites to store consent.

Upon completing a comprehensive implementation and gathering extensive feedback from the research community, we plan to engage with standards organizations like W3C. The browser-level cookie categorization and consent management approach could be proposed as a web standard, complementing existing privacy-focused initiatives. We believe standardization would be a crucial step toward widespread adoption by browser vendors and eventual integration into the web ecosystem, ultimately enhancing user privacy protection across the internet.

Acknowledgment

The authors would like to thank the anonymous reviewers and the shepherd for their excellent feedback on the paper. This work is supported in part by funding from the Prime Minister's Research Fellowship and Google.

References

[1] "Dnt - http — mdn," https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/DNT, (Accessed on 10/11/2024).

[2] "Global privacy control — take control of your privacy," https://globalprivacycontrol.org/, (Accessed on 10/10/2024).

[3] "Iab tech lab," https://iabtechlab.com/standards/ccpa/, (Accessed on 10/11/2024).

[4] "Icc uk cookie guide - cookie-cat," [Online; accessed 2025-02-12]. [Online]. Available: https://cookie-cat.co.uk/about/icc-uk-cookie-guide/

[5] "Tracking the trackers: Early results," https://cyberlaw.stanford.edu/blog/2011/07/tracking-trackers-early-results/, (Accessed on 10/11/2024).

[6] "Cookies, the gdpr, and the eprivacy directive - gdpr.eu," 5 2019, [Online; accessed 2025-02-12]. [Online]. Available: https://gdpr.eu/cookies/

[7] "CookieBlock & CookieAudit: Fixing cookie consent with ML." Boston, MA: USENIX Association, Aug. 2022.

[8] "Iab europe transparency & consent framework policies - iab europe," https://iabeurope.eu/iab-europe-transparency-consent-framework-policies/, 2023.

[9] J. A. Alharbi, A. S. Albesher, and H. A. Wahsheh, "An empirical analysis of e-governments' cookie interfaces in 50 countries," *Sustainability*, vol. 15, no. 2, p. 1231, 2023.

[10] N. Bielova, L. Litvine, A. Nguyen, M. Chammat, V. Toubiana, and E. Harry, "The effect of design patterns on (present and future) cookie consent decisions," in *USENIX Security Symposium. USENIX Association*, 2024.

[11] D. Bollinger, K. Kubicek, C. Cotrini, and D. Basin, “Automating cookie consent and GDPR violation detection,” in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 2893–2910. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/bollinger

[12] C. Bosch, B. Erb, F. Kargl, H. Kopp, and S. Pfattheicher, “Tales from the dark side: Privacy dark strategies and privacy dark patterns,” *Proceedings on Privacy Enhancing Technologies*, 2016.

[13] A. Bouhoula, K. Kubicek, A. Zac, C. Cotrini, and D. Basin, “Automated large-scale analysis of cookie notice compliance,” in *33st USENIX Security Symposium (USENIX Security 24)*. Philadelphia, PA: USENIX Association, Aug. 2024, p. TBA. [Online]. Available: https://www.usenix.org/conference/usenixsecurity24/presentation/bouhoula

[14] C. M. Bruhner, D. Hasselquist, and N. Carlsson, “Bridging the privacy gap: Enhanced user consent mechanisms on the web,” in *Proc. NDSS Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb@ NDSS), Mar. 2023.*, 2023.

[15] California State Legislature, “California consumer privacy act of 2018,” 2018. [Online]. Available: https://oag.ca.gov/privacy/ccpa

[16] S. Calzavara, G. Tolomei, A. Casini, M. Bugliesi, and S. Orlando, “A supervised learning approach to protect client authentication on the web,” *ACM Trans. Web*, vol. 9, no. 3, jun 2015. [Online]. Available: https://doi.org/10.1145/2754933

[17] G. Chrome, “Consent-o-matic,” https://chromewebstore.google.com/detail/consent-o-matic/mdjildafknihdffpkfmmpnpoiajfjnjd?hl=en, 2020, (Accessed on 06/17/2024).

[18] L. F. Cranor, M. Arjula, and P. Guduru, “Use of a p3p user agent by early adopters,” ser. WPES ’02. New York, NY, USA: Association for Computing Machinery, 2002, p. 1–10. [Online]. Available: https://doi.org/10.1145/644527.644528

[19] M. Degeling, C. Utz, C. Lentzsch, H. Hosseini, F. Schaub, and T. Holz, “We Value Your Privacy... Now take some cookies: Measuring the GDPR’s impact on web privacy,” in *Proceedings of the 26th Network and Distributed System Security Symposium (NDSS ’19)*, 2019.

[20] N. Demir, T. Urban, N. Pohlmann, and C. Wressnegger, “A large-scale study of cookie banner interaction tools and their impact on users’ privacy,” *Proceedings on Privacy Enhancing Technologies*, 2024.

[21] “Guidance from the Conference of Independent Data Protection Supervisory Authorities of the Federal Government and the States of 20 December 2021 (OH Telemedia 2021, V.1.1),” https://www.datenschutzkonferenz-online.de/media/oh/20211220_oh_telemedien.pdf, accessed on 2024.09.03.

[22] “Explanation of the standard of the AP on the withdrawal of consent for cookie banners 01 March 2024,” https://www.autoriteitpersoonsgegevens.nl/documenten/normuitleg-ap-over-intrekken-van-toestemming-bij-cookiebanners, accessed on 2024.09.03.

[23] European Data Protection Board, “Guidelines 3/2022 on dark patterns in social media platform interfaces: How to recognise and avoid them,” Mar. 2022. [Online]. Available: https://edpb.europa.eu/our-work-tools/documents/public-consultations/2022/guidelines-32022-dark-patterns-social-media_en

[24] European Parliament and Council of the European Union, “Regulation (EU) 2016/679 of the European Parliament and of the Council,” 2016. [Online]. Available: https://data.europa.eu/eli/reg/2016/679/oj

[25] C. M. Gray, C. Santos, N. Bielova, M. Toth, and D. Clifford, “Dark patterns and the legal requirements of consent banners: An interaction criticism perspective,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–18.

[26] H. Habib, M. Li, E. Young, and L. Cranor, ““okay, whatever”: An evaluation of cookie consent interfaces,” in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: https://doi.org/10.1145/3491102.3501985

[27] P. Hausner and M. Gertz, “Dark patterns in the interaction with cookie banners,” *arXiv preprint arXiv:2103.14956*, 2021.

[28] X. Hu, N. Sastry, and M. Mondal, “Cccc: Corralling cookies into categories with cookiemonster,” in *Proceedings of the 13th ACM Web Science Conference 2021*, ser. WebSci ’21. New York, NY, USA: Association for Computing Machinery, 2021, p. 234–242. [Online]. Available: https://doi.org/10.1145/3447535.3462509

[29] S. Human, “Advanced data protection control (adpc): An interdisciplinary overview,” *arXiv preprint arXiv:2209.09724*, 2022.

[30] IAB, “Index - global vendor list,” https://register.consensu.org/Translation, (Accessed on 09/01/2024).

[31] Information Commissioner’s Office (ICO), “Call for views on “consent or pay” business models,” 2020. [Online]. Available: https://ico.org.uk/cookies-call-for-views-202403

[32] G. P. Kancherla, N. Bielova, C. Santos, and A. Bichhawat, “Measuring compliance of consent revocation on the web,” *arXiv preprint arXiv:2411.15414*, 2024.

[33] G. P. Kancherla, D. Goel, and A. Bichhawat, “Least privilege access for persistent storage mechanisms in web browsers,” *arXiv preprint arXiv:2411.15416*, 2024.

[34] D. Klein, M. Musch, T. Barber, M. Kopmann, and M. Johns, “Accept all exploits: exploring the security impact of cookie banners,” in *Proceedings of the 38th Annual Computer Security Applications Conference*, 2022, pp. 911–922.

[35] S. Macbeth, “Autoconsent – get this extension for firefox (en-us),” https://addons.mozilla.org/en-US/firefox/addon/autoconsent/, 2020, (Accessed on 06/17/2024).

[36] M. Nouwens, I. Liccardi, M. Veale, D. R. Karger, and L. Kagal, “Dark patterns after the GDPR: scraping consent pop-ups and demonstrating their influence,” *CoRR*, vol. abs/2001.02479, 2020. [Online]. Available: http://arxiv.org/abs/2001.02479

[37] Å. A. Nyre, K. Bernsmed, S. Bo, and S. Pedersen, “A server-side approach to privacy policy matching,” in *2011 Sixth International Conference on Availability, Reliability and Security*. IEEE, 2011, pp. 609–614.

[38] OneTrust, “All you need to know about cookies — cookiepedia,” https://cookiepedia.co.uk/, 2020.

[39] F. Roesner, T. Kohno, and D. Wetherall, “Detecting and defending against Third-Party tracking on the web,” in *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. San Jose, CA: USENIX Association, Apr. 2012, pp. 155–168. [Online]. Available: https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/roesner

[40] M. H. Ronald Leenes, Jan Schallaböck, “Prime - privacy and identity management for europe — portal for the prime project,” https://prime-project.eu/, 2006.

[41] I. Sanchez-Rola, M. Dell’Amico, P. Kotzias, D. Balzarotti, L. Bilge, P.-A. Vervier, and I. Santos, “Can i opt out yet? gdpr and the global illusion of cookie control,” in *Proceedings of the 2019 ACM Asia conference on computer and communications security*, 2019, pp. 340–351.

[42] C. Santos, N. Bielova, and C. Matte, “Are cookie banners indeed compliant with the law?” *Technology and Regulation*, vol. 2020, pp. 91–135, Dec. 2020. [Online]. Available: https://inria.hal.science/hal-02875447

[43] C. Santos, A. Rossi, L. Sanchez Chamorro, K. Bongard-Blanchy, and R. Abu-Salma, “Cookie banners, what’s the purpose? analyzing cookie banner text through a legal lens,” in *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, ser. WPES ’21. New York, NY, USA: Association for Computing Machinery, 2021, p. 187–194. [Online]. Available: https://doi.org/10.1145/3463676.3485611

[44] M. Toth, N. Bielova, and V. Roca, “On dark patterns and manipulation of website publishers by CMPs,” *Proceedings on Privacy Enhancing Technologies (PoPETs)*, vol. 2022, no. 3, pp. 478–497, 2022. [Online]. Available: https://inria.hal.science/hal-03577024

[45] C. Utz, M. Degeling, S. Fahl, F. Schaub, and T. Holz, “(un) informed consent: Studying gdpr consent notices in the field,” in *Proceedings of the 2019 acm sigsac conference on computer and communications security*, 2019, pp. 973–990.

[46] W3C, “The platform for privacy preferences 1.1 (p3p1.1) specification,” https://www.w3.org/TR/P3P11/, 2006, (Accessed on 06/17/2024).

[47] S. Zimmeck, O. Wang, K. Alicki, J. Wang, and S. Eng, “Usability and enforceability of global privacy control,” *Proceedings on Privacy Enhancing Technologies*, vol. 2023, no. 2, 2023.