

Usability and Enforceability of Global Privacy Control

Sebastian Zimmeck*
szimmeck@wesleyan.edu
Wesleyan University
United States

Oliver Wang
owang@wesleyan.edu
Wesleyan University
United States

Kuba Alicki†
kuba.alicki@princeton.edu
Princeton University
United States

Jocelyn Wang
jwang06@wesleyan.edu
Wesleyan University
United States

Sophie Eng
seng@wesleyan.edu
Wesleyan University
United States

ABSTRACT

Web tracking by ad networks and other data-driven businesses is often privacy-invasive. Privacy laws, such as the California Consumer Privacy Act, aim to give people more control over their data. In particular, they provide a right to opt out from web tracking via privacy preference signals, notably Global Privacy Control (GPC). GPC holds the promise of enabling people to exercise their opt out rights on the web. Broad adoption of GPC hinges on its usability. In a usability survey we find that 94% of the participants would turn on GPC indicating a need for such efficient and effective opt out mechanism. 81% of the participants in our survey also have a correct understanding of what GPC does ensuring that their intent is accurately represented by their choice.

The effectiveness of GPC is dependent on whether websites' GPC compliance can be enforced. A site's GPC compliance can be analyzed based on privacy flags, such as the US Privacy String, which is used on many sites to indicate the opt out status of a web user. Leveraging the US Privacy String for GPC purposes we implement a proof-of-concept browser extension that successfully and correctly analyzes sites' GPC compliance at a rate of 89%. We further implement a web crawler for our browser extension demonstrating that our analysis approach is scalable. We find that many sites do not respect GPC opt out signals despite being legally obligated to do so. Only 54/464 (12%) sites with a US Privacy String opt out users after having received a GPC signal.

KEYWORDS

Global Privacy Control, GPC, Do Not Sell, Do Not Share, Do Not Track, DNT, Opt Out, Privacy Preference Signals, Privacy Choice, Privacy Rights, CCPA, CPRA, Privacy Law, Usability, Web Privacy

1 INTRODUCTION

People lack control over their data on the web. Many websites track people's browsing habits, geolocations, or other personal information. Such tracking is often not transparent and difficult to control.

*Corresponding author.

†Work performed while at Wesleyan University.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2023(2), 265–281

© 2023 Copyright held by the owner/author(s).
<https://doi.org/10.56553/poops-2023-0052>

Lawmakers are increasingly enacting privacy laws to protect people's privacy. Such laws are necessary to establish a baseline of privacy rights. The European Union modernized privacy law for the web with the General Data Protection Regulation (GDPR) [24]. Stateside, after a fairly narrow Nevada privacy law [54], California adopted the more comprehensive California Consumer Privacy Act (CCPA) [14]. The CCPA provides California consumers with data access, deletion, and opt out rights. Those rights are being further extended by the California Privacy Rights Act (CPRA) [14]. Virginia [28], Colorado [27], Utah [75], and Connecticut [74] followed suit with their own privacy laws. Privacy laws as such do not directly lead to more privacy protection. The privacy rights provided by these laws need to be enforced. Exercising a privacy right — opting out from web tracking, accessing data, or deleting it — should be just as simple as the process of collecting or sharing data. The GDPR makes this principle explicit: any withdrawal of consent must be as easy as giving it.¹

The opt out right is especially important because it controls whether or not people's data enters the online ad ecosystem. There will be no need to access or delete data if it was not stored. Different from access and deletion rights, opt out rights are not subject to the risk of unauthorized access or deletion and, thus, easier to implement for website operators. However, current opt out implementations are ineffective and inefficient. Opting out site-by-site via Do Not Sell links is effectively unusable [50]. While laws increasingly provide opt out rights, they will not matter much if they cannot be practically exercised. The CCPA specifies the opt out right in detail. Per the CCPA, California consumers can direct businesses to not sell their personal information to third parties.² "Selling" is defined as obtaining any monetary gain, for example, by disclosing a consumers' personal information to an ad network on a website via third party cookies. The critical centerpiece for making the opt out right effectively and efficiently usable are privacy preference signals [35].

Privacy preference signals enable web users, under the CCPA and other laws, to send Do Not Sell and Do Not Share signals to websites via browsers or browser extensions. The state laws and regulations in California, Colorado, and Connecticut provide provisions for privacy preference signals. The California Consumer Privacy Act Regulations (CCPA Regulations) specify that "[i]f a business collects personal information from consumers online, the business shall treat user-enabled *global privacy controls*, such as a

¹GDPR Art. 7(3).

²California Civil Code §1798.120(a).

browser plug-in or privacy setting, device setting, or other mechanism, that communicate or signal the consumer’s choice to opt out of the sale of their personal information as a valid request [...].”³ Similarly, the Colorado Privacy Act (CPA) allows consumers to opt out via “a user-selected universal opt-out mechanism” and the Connecticut Data Privacy Act (CDPA) via “an opt-out preference signal.”⁴ Per the GDPR, a “data subject may exercise his or her right to object by automated means using technical specifications.”⁵ In this study we focus on California law as it provides broad recognition for privacy preferences signals and already began enforcing Global Privacy Control (GPC) [72].

GPC is a privacy preference signal developed by a coalition of privacy organizations, publishers, browser vendors, extension developers, and academics that aims to make it easier for people to exercise their Do Not Sell and Do Not Share rights [29]. In various ways GPC is similar to Do Not Track (DNT) [78]. Just as DNT, GPC is a binary signal; an on-off switch for web tracking. However, different from DNT, people can turn on GPC for all sites they visit or only for some if they so desire. Also, while DNT is purely header-based, GPC can be implemented via the Sec-GPC request header or the GPC JavaScript DOM property [30]. The latter usually enables sites much quicker to determine the opt out status of a site visitor. Most importantly, GPC is required by law and enforced by the Office of the California Attorney General (OAG) [71–73]. To propagate consumers’ opt outs of sale through the ad ecosystem the Interactive Advertising Bureau (IAB) developed the US Privacy String [39, 40]. Upon receiving a GPC signal site operators would need to change the US Privacy String to reflect the opt out. In this study we evaluate how people can effectively and efficiently opt out via GPC.

- (1) In two online usability surveys, we evaluate whether study participants understand what GPC does and the extent to which they would make use of it. As 94% of participants would enable GPC and 81% correctly understand that it limits data selling and sharing, our results suggest a real-world need for GPC as well as that its functionalities can be conveyed comprehensibly. (§3)
- (2) We describe a methodology leveraging the US Privacy String for detecting whether a site respects GPC and implementing it in a browser extension as an automated compliance tool. Our implementation performs an accurate GPC compliance analysis at a rate of 89% on a test set of 100 sites.⁶ (§4)
- (3) We perform a GPC compliance enforcement analysis of 64 top websites and find that only 27 sites are compliant. Further, using our GPC extension in combination with a web crawler we analyze 464 sites for GPC Compliance. Our results show that the current status of GPC compliance is low with only 54/464 (12%) sites respecting GPC. (§5)

2 BACKGROUND AND RELATED WORK

Since its beginnings in the mid-1990s, online advertising has become more intricate and over time evolved into a complex system

³CCPA Regulations §999.315(c) (emphasis added).

⁴CPA §6-1-1306(1)(a)(IV)(B), §6-1-1313, CDPA §6(e)(1)(A)(ii).

⁵GDPR Art. 21(5). While the GDPR is based on the opt in principle, it also contains an “opt out” as people have the right to withdraw consent for processing data or raise an objection to such processing, e.g., per GDPR Art. 6, 7, and 21.

⁶Our code is available at <https://github.com/privacy-tech-lab/gpc-optmeowt>.

comprised of publishers, ad networks, data brokers, and ad tech companies [49, 82]. At the center of many publishers’ business models are personalized ads that are targeted to people based on their web browsing activity and other data collected about them.

2.1 Web Tracking and Ad Blocking

To target people with personalized ads, publishers integrate third party software components, notably, those of ad networks, into their sites. However, the presence and behaviors of those components are not apparent to the average site visitors. People are often not aware that they are being tracked, or if they are, they do not know who is tracking them or which types of data the trackers receive. Most people also do not know how to effectively identify or control behavioral profiling and other privacy-invasive practices [66]. However, an increasing number of people are using ad blockers. A recent survey covering web users from the Netherlands showed that ad blockers appear to be used at least occasionally by 30% of the respondents [9]. While some publishers block ad blockers creating an ad blocking arms race [55, 84], many ad networks and publishers seem motivated to evolve their business models towards higher privacy protections due to an increasingly skeptical audience and regulatory scrutiny [26]. This trend is further amplified by browser vendors moving ad blocking features directly into their browsers.

2.2 Ad Industry Third Party Opt Out Tools

Different from ad blockers, industry opt out tools appear to be far less used with only 7% of people reporting occasionally using those [9]. The IAB, Digital Advertising Alliance (DAA), and Network Advertising Initiative (NAI) provide self-serve opt out sites, such as the YourAdChoices site [23] or CCPA opt out sites [22, 38], which, however, are unsuitable as CCPA opt out mechanisms [45]. Those sites usually work by recording people’s opt out choices in third party cookies accessible to the affected ad networks. However, this mechanism is impractical for various reasons. First users would need to repeat their choices for each of their browsers. Further, it is doubtful that many people would find their way to the sites with the choice tools. The identification of privacy choice links in privacy policy text and surfacing the links to the user may ameliorate this problem [5, 64], though, it is even easier to write the opt out cookies directly into the browser storage via a browser extension [87]. In any event, as browser vendors are phasing out support for third party cookies [11], cookie-based ad industry opt out sites will be of little relevance going forward. In the iOS ecosystem, Apple’s App Tracking Transparency framework [3] is much more substantial than industry efforts on the web. Google is planning a similar effort for Android [15].

2.3 Opting Out Directly on First Party Sites

The CCPA Regulations specify that a business that sells personal information shall post a “Do Not Sell My Personal Information” link on its homepage [70].⁷ However, opting out by web form, the most common CCPA Do Not Sell opt out mechanism on the web, is inefficient by design because it requires consumers to opt out

⁷CCPA Regulations §999.306(b)(1).

for every business individually [50]. Many opt out implementations also deter people from opting out by nudging them towards staying opted in [56]. Some businesses have adopted practices that are onerous, confusing, or even non-compliant with the law, e.g., by mandating unnecessary identity verification contrary to what the CCPA Regulations require [50].⁸ People should never have to input information that a business does not already have, and opt out mechanism design should be kept simple, for example, by not forcing people to opt out for different purposes separately [56]. Numerous sites also have missing, incorrectly placed, incorrectly worded, or broken Do Not Sell links highlighting the importance of a strong enforcement regime [56]. Overall, the large number of websites collecting data makes the adoption of Do Not Sell browser signals or future legislation to limit data sales critically important for enhancing people's privacy at scale [56].

2.4 P3P and DNT

The first major privacy preference signal implementation, the Platform for Privacy Preferences Project (P3P) [17, 18, 21], enabled people to delegate privacy decisions to user agents that could automatically react to the privacy practices on websites according to their machine-readable privacy policies. However, at the time mainstream adoption of P3P was hindered due to a rather privacy-hostile environment [48] and natural language privacy policies already having been established as de facto privacy policy standard [88]. However, P3P sparked a flurry of innovations, e.g., in privacy preference languages [2, 20], user interface design for privacy agents [19, 62], privacy "nutrition" labels [46], or P3P-enabled web search [12]. While the use of P3P policies in today's web environment would subject users to browser fingerprinting risks, at the time its development was an important milestone in the development of privacy preference signals. Two important lessons from P3P — keeping privacy preference signals as simple as possible and working directly with browser vendors on their implementation [65] — were taken on by the Tracking Protection Working Group [79] in the development of DNT [78]. DNT was designed as a binary signal for people to express their opt out from tracking per the California Online Privacy Protection Act (CalOPPA) [13]. However, a study at the time showed that only two out of hundreds of ad networks respected DNT [4]. The reason for this lack of adoption was that CalOPPA does not require recipients of DNT signals to respect those but only to *disclose* whether they do so.⁹ Consequently, most sites simply disclose in their privacy policies that they do not respect DNT.

2.5 GPC

GPC can be considered a spiritual successor to DNT [83]. GPC is a privacy preference signal developed by a coalition of privacy organizations, publishers, browser vendors, extension developers, and academics that aims to make it easier for people to exercise their opt out rights [29].

2.5.1 Legal Bindings. GPC's initial use case is the CCPA, but it can also be applied to the GDPR and other laws. The proposed

GPC standard does not prescribe any legal meaning to a GPC signal but rather leaves this determination to local lawmakers and regulators [30]. Thus, a GPC signal may have different legal bindings in different jurisdictions. The Office of the California Attorney General is interpreting and enforcing GPC signals as Do Not Sell request [6]. GPC can be applied to the GDPR as well. Under the GDPR sending a GPC signal could be interpreted as the withdrawal of consent and objection to the processing of personal data by third parties [8].¹⁰ However, challenges of applying GPC to the GDPR are the need for uniform interpretation by the various European Data Protection Authorities and their agreement to enforce it [58].

2.5.2 Implementation and Standardization. GPC is a binary signal. People can turn on GPC for all sites they visit or only for some. Most people make binary decisions even when more choices are available [59, 76]. Browser and extension vendors can implement the HTTP Sec-GPC request header or the GPC JavaScript DOM property [30]. Brave, Mozilla, DuckDuckGo, and Disconnect are some of the current implementers [29]. Websites receiving GPC signals do not need to keep the state of a user's opt out status as the signal will be included in every request. The GPC group provides sample reference implementations for detecting GPC signals [31]. Publishers supporting GPC are, among others, The New York Times, The Washington Post, and Automattic/WordPress. Consent Management Platforms (CMPs), such as OneTrust [57], also integrate GPC into their products thereby enabling websites relying on CMPs to use GPC. As GPC signals are available with the first web request, publishers and ad networks are immediately aware of it and can determine how to serve ads, e.g., targeted or contextual. GPC will be included as an extension in the IAB's Open Real-Time Bidding (OpenRTB) protocol [41]. The implementation of GPC goes hand in hand with its standardization. GPC is currently a W3C draft standard [86] discussed in the W3C Privacy Community Group [81]. It is not enough to tag on privacy via cookie banners or similar measures, privacy must become part of the fabric of the Internet. In principle, GPC can be applied beyond the web, for example, on mobile and IoT devices.

2.5.3 Enforcement. The DNT experience shows that the enforcement of privacy preference signals is critical. If they would not be enforced, it is likely that most website operators would ignore privacy preference signals. Accordingly, the CCPA Regulations mandate that recipients of Do Not Sell signals respect such as valid opt outs: "If a business collects personal information from consumers online, the business shall treat user-enabled global privacy controls, such as a browser plug-in or privacy setting, device setting, or other mechanism, that communicate or signal the consumer's choice to opt-out of the sale of their personal information as a valid request [...]."¹¹ The privacy laws of Colorado and Connecticut also allow opt outs via universal opt-out mechanisms and opt-out preference signals, respectively.¹² Per the OAG, "[u]nder law, [GPC] must be honored by covered businesses as a valid consumer request to stop the sale of personal information" [73]. The OAG recently enforced GPC against a business that did not respect consumers' opt outs via GPC [72]. Moreover, the CPRA is tasking the newly established

⁸CCPA Regulations §999.315(g).

⁹California Business and Professions Code §22575(b)(5).

¹⁰GDPR Art. 7(3), 21(5).

¹¹CCPA Regulations §999.315(c).

¹²CPA §6-1-1306(1)(a)(IV)(B), §6-1-1313, CDPA §6(e)(1)(A)(ii).

California Privacy Protection Agency to adopt regulations to define the requirements and technical specifications for such opt outs.¹³ Site operators can show site visitors that they comply with GPC by including a `/.well-known/gpc.json` resource on their site. In addition, increasingly methods are being developed to check sites' CCPA compliance. For example, website ranking and lead generation services can be used to determine if a site integrates third party advertising or analytics components, i.e., sells data, and if the site's number of unique visitors from California is above the threshold for CCPA applicability [77].

2.5.4 GPC as Default Setting. GPC settings should accurately represent people's privacy preferences.¹⁴ Whether a default setting in a browser or browser extension can represent a user's preference accurately may differ by jurisdiction. In some jurisdictions, particularly, those with opt in laws, turning on GPC by default upon browser installation is plausible while laws in opt out jurisdictions may require an affirmative act for GPC signals to be considered valid expressions of intent to opt out. However, such affirmative act may not necessarily need to be explicit. For example, per the CCPA, a consumer's choice of using privacy-preserving browsers or other tools is considered a sufficiently deliberate act that is interpreted as a consumer expression of a preference to not have personal information sold or shared. "The consumer exercises their choice by affirmatively choosing the privacy control [...] including when utilizing privacy-by-design products or services" [69].

2.5.5 Identification and Authentication. As GPC signals on their own do not contain any identifying information, the identification and authentication of individuals in connection with their privacy preference expressions depends on external mechanisms. For example, if a user is logged in on a website with their email address, sending GPC signals would allow the site to correlate the signals to the email address and apply the opt out to all browsers and devices associated with the same address.¹⁵ On the other hand, if a site without authentication functionality identifies people based on user IDs, ad IDs, or other pseudonyms, those can be used as identifiers. In general, by whichever mechanism a user is identified on a site, the same mechanism can also be used for opt out identification. For example, for a user identified by a cookie ID it would be unnecessary and, in fact, a dark pattern, to ask for their email address for the opt out. The situation is not different from current cookie-based opt out mechanisms (§2.2 and §2.3). Those do not rely on identification information beyond a cookie ID either. While GPC does not use cookies but HTTP headers and the DOM instead, this technological difference does not make a difference for identification purposes.

2.5.6 Browser Fingerprinting. Adding GPC to a browser can increase its fingerprinting surface. As GPC is either turned on or off for a given site, attackers will have one additional bit available to fingerprint a user's browser configuration on that site. To avoid this increase Brave enables GPC by default for all its users in an unconfigurable setting [67]. For freely configurable GPC settings

¹³California Civil Code §1798.185(a)(19)(A), 1798.185(d). The German Telemediengesetz has a similar provision for adoption of a technical specification in §26.

¹⁴California Civil Code 1798.185(a)(19)(A)(iii), CPA §6-1-1313(2)(c), CDPA §6(e)(1)(A)(ii)(II).

¹⁵See CCPA Regulations §999.315(c).

the fingerprinting risk depends on the extent to which people turn GPC on and off for different sites and how many of those sites an attacker can observe. However, overall GPC represents minor risks given its binary states [36].

2.6 ADPC and DRP

The Advanced Data Protection Control (ADPC) – similar to P3P and focused on enabling cookie consent and other privacy choices under the GDPR and ePrivacy Directive – establishes bidirectional communication between websites and users [1, 36]. ADPC is discussed in the W3C Consent Community Group [80]. The Data Rights Protocol (DRP) is a technical standard for exchanging data rights requests under the CCPA [16]. While privacy preference signals are focused on people's opt out choices, many privacy laws also contain rights for users to access or delete data, among others. Those rights are more difficult to implement electronically because of the necessary authentication to ensure that a user is allowed to have access to or delete the requested data. DRP aims to address this challenge by defining a protocol with a set of standardized request and response endpoints between businesses. With increasing rates of privacy preference signal adoption and emerging new signals a likely problem will be that users send conflicting signals [34].

2.7 Downstream Propagation of Opt Outs

Usually, when people opt out, ad networks and other third parties in the online ad ecosystem will still receive the covered data [85]. It is just that those third parties are not allowed to use it anymore and should not store it. Thus, for any opt out – whether by opt out cookies, privacy preference signals, or another mechanism – an individual's choice needs to be attached to any downstream web request passing along data the choice relates to. To propagate an opt out choice accordingly per the GDPR the IAB Europe introduced the Transparency & Consent Framework (TCF) [42]. The TCF's Transparency & Consent String (TC String) is intended to transmit a user's consent for processing personal data for certain purposes and vendors as well as whether a vendor is allowed to process the user's data based on legitimate interest [43]. However, in 2022 the Belgian Data Protection Authority ruled that the TCF does not comply with the GDPR and would need to be revised [7].

In addition to the TCF, the IAB also introduced the CCPA Compliance Framework for Publishers & Technology Companies [38]. Its US Privacy String is a four-character string that identifies if a consumer has opted out from the sale of personal information under the CCPA [39, 40]. The TCF and CCPA Compliance Framework are part of IAB's Global Privacy Platform, which is intended to comprehensively transmit consent and choice signals from sites and apps to ad networks [44]. Google and Meta introduced their own privacy flags for CCPA compliance. If turned on by a site, Google's Restricted Data Processing (RDP) privacy flag signals to Google not to sell or share data with third parties [32]. Upon receiving a GPC signal, sites may choose to enable RDP [32]. Meta's Limited Data Use (LDU) privacy flag works similarly [25].

3 USABILITY OF GPC

Interacting with Do Not Sell mechanisms on individual websites is burdensome and ineffective [50]. As some businesses have adopted

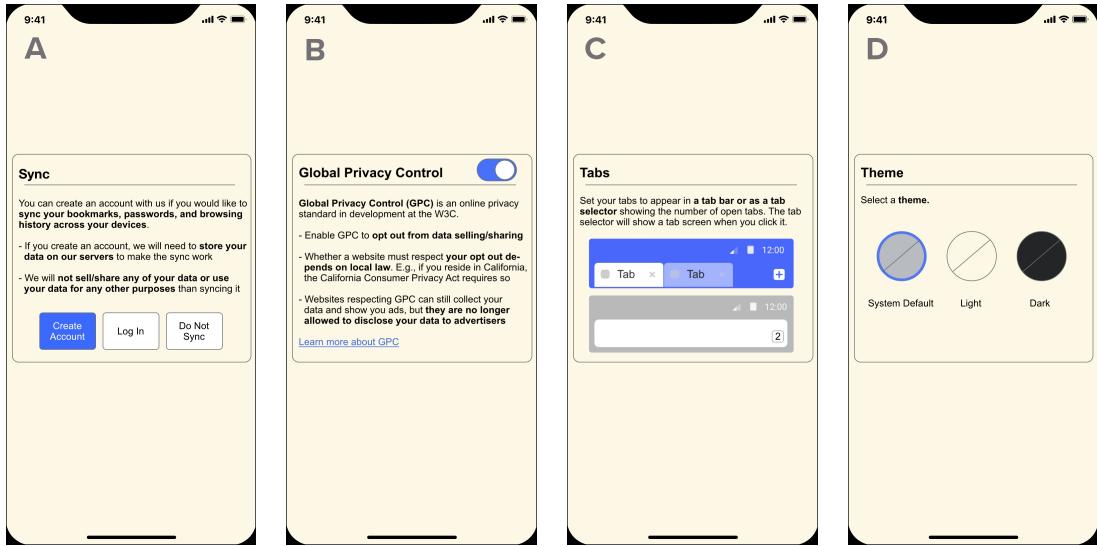


Figure 1: Screenshots of the GIFs for the four mobile browser features presented in the Browser Setup Survey. The blue highlighting of the available options alternated so that none of the options appeared as default. The explanation in screenshot B was also used in the GPC Survey (Figure A.1 in the Appendix).

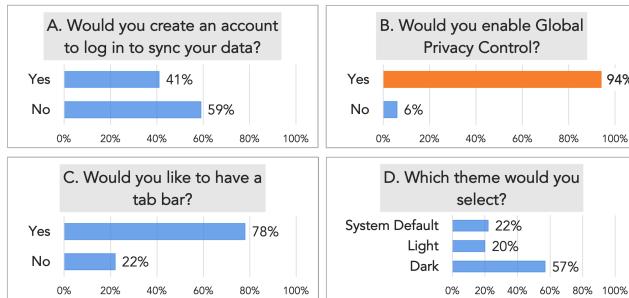


Figure 2: Results of our Browser Setup Survey. The survey questions and answer percentages relate to the screenshots shown in Figure 1. 94% of the participants in the Browser Setup Survey opted to enable GPC, which is a statistically significant outcome.

practices that are onerous, confusing, or even non-compliant with the law [50], GPC holds the promise of enabling web users to make use of their opt out rights more efficiently and effectively in a manner sanctionable by applicable laws. The CCPA Regulations require businesses to honor GPC signals as valid opt out choices.¹⁶ So does recent enforcement practice by the OAG [71–73]. While GPC is not yet broadly available in all browsers and few know about it, we performed two online surveys to explore if people would enable it (§3.2) and if they would understand what it does (§3.3).

3.1 Methodology

We obtained IRB exemption declarations for our surveys from our institution and recruited 50 participants for our first survey (Browser Setup Survey) and 50 for our second survey (GPC Survey) on the crowdsourcing platform Prolific [61]. We selected a sample

size of 50 participants for each survey. We excluded 3/100 participants who failed to answer an attention test question correctly leaving us with 49 participants in the Browser Setup Survey and 48 in the GPC Survey. We paid each participant as recommended by Prolific \$0.80 for the Browser Setup Survey, which took participants a median time of 2 minutes to complete, and \$1.60 for the GPC Survey, which took participants a median of 5 minutes and 37 seconds to complete. In our consent forms we let participants know that they are free to decline to participate, to end participation at any time for any reason, or to refuse to answer any individual question. All participants completed the full surveys and received the full compensation.

For both surveys we screened participants to be at least 18 years old, California residents, fluent in English, with at least 50 tasks previously performed on Prolific, and 100% acceptance rate for their previous tasks. While we cannot exclude it, we are not aware that any study participant knew about GPC beforehand. Given that Brave and Firefox – the only browsers implementing GPC, for the time being – do not provide dedicated user interfaces for GPC and do not explicitly provide in-browser alerts about this new functionality, many people do not yet know about GPC. Participants could only take part in one of our surveys. The Browser Setup Survey consisted of multiple choice questions and the GPC Survey of both free-form and multiple choice questions.¹⁷ For multiple choice questions that did not depend on any order, we shuffled the answer choices. Across both surveys 8% of participants were 18 or 19 years old, 46% between 20 and 29, 24% between 30 and 39, and 21% were 40 or older. 56% were male and 44% female. All demographic data was automatically provided to us by Prolific according to

¹⁶CCPA Regulations §999.315(c).

¹⁷The complete survey questionnaires are shown in Appendices A.1 and A.2.

their default procedure.¹⁸ The set of survey participants is not a representative sample of the population using the web.

3.2 Would People Enable GPC?

Brave and Firefox do not provide a dedicated user interface for GPC. Brave enables GPC by default for all users and the setting is unconfigurable [67]. Firefox disables GPC by default and provides an option to enable GPC via the general about:config settings [53]. Neither browser notifies users about GPC upon installation. Thus, we explore in our Browser Setup Survey whether users would enable GPC when presented as part of a browser setup process. To avoid biasing the participants we described our survey on Prolific as research on web browser features without specifying our particular focus on privacy or GPC.¹⁹ We asked participants to make setup choices for four features in a mobile browser: (A) account functionality to sync bookmarks and other data, (B) GPC, (C) a tab bar, and (D) dark and light themes (Figure 1). For the non-GPC settings, the features and their presentation are loosely based on those provided by the mobile Vivaldi browser. For the GPC setting, we showed participants a screenshot for enabling GPC with an explanation of what it does as shown in Figure 1B. The Browser Setup Survey mirrors a real browser setup process. The work of both tasks, i.e., selecting browser interface features via buttons and switches, is the same. We aimed to minimize the privacy paradox – the notion that people act differently from how they say they would act when it comes to privacy decisions.²⁰ Thus, we presented the experimental setup closer to a real browser setup task that is more behavioral in nature rather than directly asking participants about enabling GPC.

46/49 (94%) of the participants in our Browser Setup Survey opted to enable GPC (Figure 2B). Using Fisher's exact test comparing this distribution against a uniform distribution we obtain a statistically significant result with $p < 0.05$. The preference for enabling GPC is the largest majority preference among all features we presented. Only 78% of participants would like to have a tab bar, 59% would not create an account to sync their data, and 57% would like to have a dark theme. Thus, GPC seems to be a relatively popular browser feature that many web users would welcome to have and use. We confirmed this result in our GPC Survey. When asked to how many websites they would send GPC signals if they could pick them individually, 89% of the participants would pick all or most while 10% would only pick some or none (Figure 3A). Consistent with these answers, when asked why they would enable GPC, only 10% of participants replied that they would not enable GPC (Figure 3B). These levels of enabling GPC are also consistent with the rates of opting out from app tracking on iOS, which, based on reports from app analytics companies, range between 60% and 95% [47]. Overall, our results suggest that people would make use of GPC if it were available to them and they would know about it.

¹⁸The demographic data includes the participants who did not answer the attention question correctly as Prolific does not allow us to associate individual data points to individual survey participants. For one participant Prolific did not provide the age.

¹⁹We disclosed to participants that we study “user preference settings in browsers,” which our IRB found to be complete and non-deceptive. Thus, it was not necessary to debrief participants.

²⁰Arguably, the privacy paradox is a myth [68]. If privacy technologies were more usable, there may be more consistency between what people say and do.

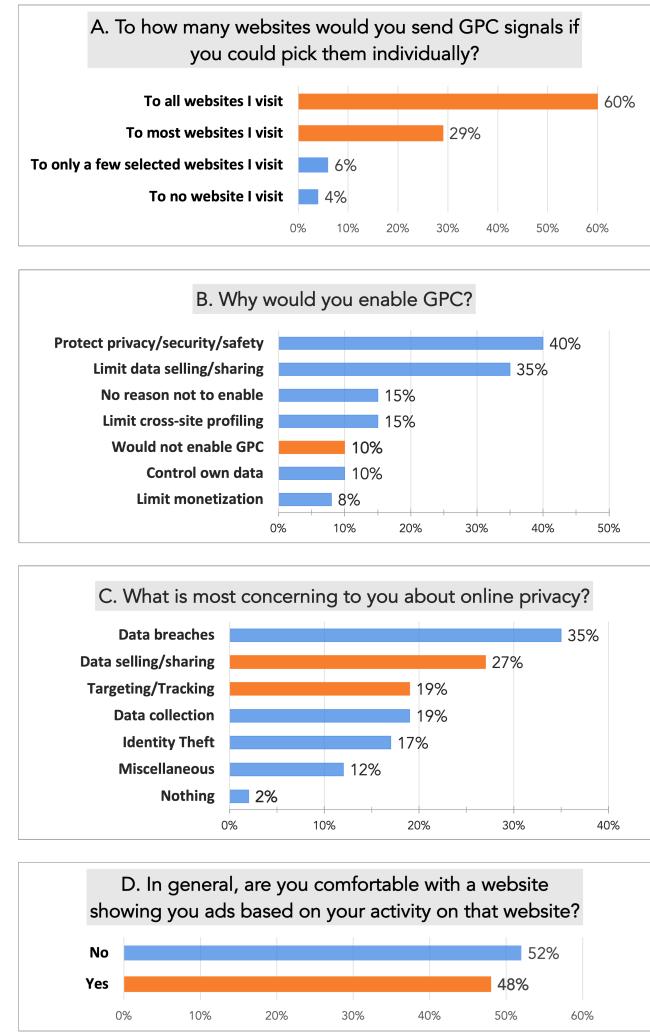


Figure 3: Results of our GPC Survey. Figures A and D show responses to multiple choice questions. Figures B and C show coded responses to free-form questions. Some participants gave multiple responses.

Since many people do not change default browser settings, the GPC default setting, which may differ depending on local law (§2.5.4), will be a critical factor for how broadly GPC signals will be disseminated across the web. In the absence of asking people or a law prescribing the default, enabling GPC is a plausible default if it represents the intent of the majority. That is the case for the 89% of participants in our GPC Survey, who would send GPC signals to all or most sites they visit. In addition, as a number of participants expressed concerns about their online privacy, especially, on data selling/sharing and targeting/tracking (Figure 3C), those concerns also justify enabling GPC by default. It should be noted, though, that the broad availability of GPC does not mean the end of all ad-based business models. Rather, GPC still leaves room for such business models to the extent they respect people's privacy. Indeed, 48% of participants in our GPC Survey expressed that they

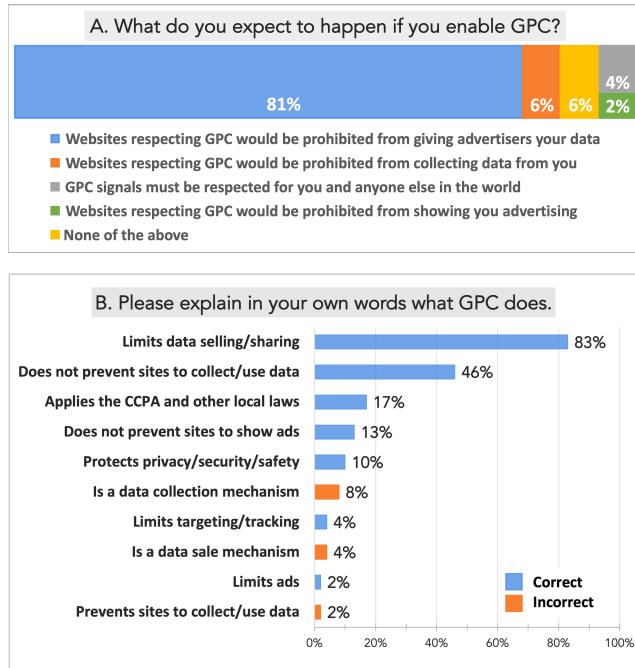


Figure 4: Results of our GPC Survey. Figure A shows what participants expect to happen if they turn on GPC (as shown in the screenshot in Figure A.1 in the Appendix). Figure B shows percentages of coded responses to the free-form question of explaining GPC. Some participants gave multiple responses.

are comfortable with receiving first party ads given that their data is not shared with advertisers or other companies (Figure 3D).

Some participants expressed doubt about GPC and disillusionment about the state of privacy on the web in general. For example, one participant was hesitant to enable GPC on the basis that GPC itself could possibly be used for tracking. Indeed, keeping the browser fingerprinting risk low is critical [67]. Others expressed that they do not care anymore as their identity had been already exposed or they may feel a false sense of security. Another participant confused GPC with DNT and thought that GPC would not do anything new. Thus, it is important to make users aware of how GPC can help them and follow through with a strong enforcement regime [77]. Of equal importance is to keep GPC convenient and easy to use as one participant explicitly was concerned about having to invest a lot of effort to improve their privacy situation. Designing *usable* GPC settings would go a long way towards GPC adoption.

3.3 Would People Understand What GPC Does?

In addition to whether people would enable GPC, our GPC Survey also provides an indication for whether they would understand what GPC does. People should enable GPC based on an informed decision that represents their intent to stop all or a certain set of websites from sharing or selling their data. Any privacy choice mechanism that does not represent a person's intent runs the risk of representing the intent of the mechanism designer or implementer. Understanding what GPC does requires at least a very basic

understanding of how the web ad ecosystem works. People should understand in broad strokes:

- (1) The difference between first and third parties
- (2) That GPC prevents selling of data to or sharing it with certain third parties
- (3) That they may still see ads
- (4) That GPC does not necessarily apply to all websites they visit, depending on local law

In addition, people should also be aware of and understand their opt out rights. As our GPC Survey participants were California residents we showed them a screenshot explaining that they have a CCPA opt out right and that GPC is mandatory in California.²¹ After showing them these explanations we asked participants in our GPC Survey what they expect would happen if they enable GPC (Figure 4A). Our results indicate that the GPC explanation we provided worked reasonably well to give participants an understanding of what GPC does. Based on the explanation, 81% of the participants in our GPC Survey correctly answered that advertisers would not receive their data if they turned on GPC.²²

To further evaluate participants' understanding of GPC we asked them in a free-form question to explain in their own words what GPC does (Figure 4B). 83% of participants explained in their answers that GPC limits data selling and sharing while 46% also included in their answer that it does not prevent sites from collecting data or using collected data themselves. Thus, these participants appear to have the correct understanding that GPC does not affect data collection practices by first parties, which was a major misunderstanding of DNT at the time [51]. However, 8% of the participants incorrectly thought that GPC itself is a data collection mechanism replacing current ad technology, and 4% thought that it is a mechanism for selling data. Given these misunderstandings, any explanation of GPC would benefit from pointing out that GPC itself is just a communications channel for privacy choices and not itself a mechanism for collecting, sharing, or selling data. The reasons participants gave for why they would enable or disable GPC (Figure 3B) lends further evidence to their correct understanding of what GPC does as the reasons were congruent with what GPC, in fact, is intended to do.

4 DETECTING GPC NON-COMPLIANCE

To the extent website operators are required by law to respect privacy choices expressed via GPC, it is necessary to have a mechanism for evaluating whether they actually do so. Otherwise, compliance could not be enforced. However, the proposed GPC standard does not provide for any compliance mechanism [30]. Thus, we describe a methodology for identifying whether certain sites that are subject to the CCPA respect GPC privacy choices (§4.1) and implement it in a browser extension as an automated compliance tool (§4.2).²³

4.1 Methodology

In the current online ad ecosystem, even if users have already opted out, many websites continue to make ad calls as usual while

²¹See Appendix, Figure A.1, which contains the same text as Figure 1B.

²²We phrased the answers to this multiple choice question differently from the language that was contained in the explanation shown in Figure 1B so that participants did not simply repeat back to us what they read without attempting to understand GPC.

²³Our code is available at <https://github.com/privacy-tech-lab/gpc-optmeowt>.

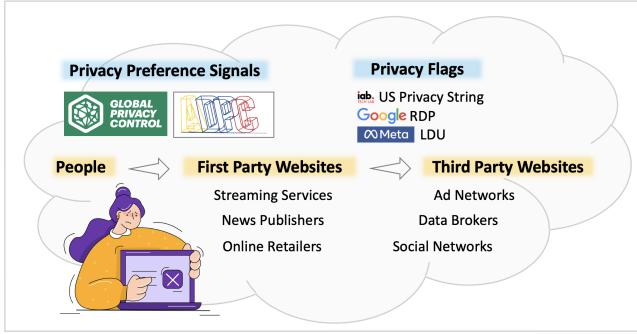


Figure 5: People make their opt out choices via privacy preference signals, such as ADPC or GPC. The signals will set a site's privacy flags, such as the US Privacy String or Google's RDP flag, to an opt out configuration. The flags are then passed to downstream ad networks and other third party sites.

signaling to downstream ad networks that users are opted out via the US Privacy String or other privacy flags [85]. Thus, it is plausible to view the changing of a privacy flag to an opt out configuration as circumstantial evidence of a site respecting a user's opt out choice.

4.1.1 Detection of Changes in the US Privacy String. The US Privacy String, which identifies if a consumer has opted out from the sale of personal information under the CCPA [39, 40], is the most commonly implemented privacy flag for passing on a user's opt out choice under the CCPA. It can be implemented in form of an HTTP header, for example, as an HTTP cookie [40], but also in form of the client-side JavaScript USPAPI [39]. In both implementations, the US Privacy String has four characters. The third character determines a user's opt out status and can have the following values: Y if a user has opted out of the sale of personal information per the CCPA, N if a user has not opted out, and “-” if the CCPA does not apply. To identify whether a site respects a user's opt out via GPC a four-pronged analysis can be applied:

- (1) Check if a site has a Do Not Sell link to determine if the site is subject to the CCPA
- (2) Check the site's US Privacy String, in particular, the third character, to determine a user's current opt out status
- (3) Send a GPC opt out signal to opt out
- (4) Check again the site's US Privacy String, in particular, the third character, to determine the user's current opt out status

4.1.2 Detection of Changes in Privacy Flags in General. In principle, this methodology can be applied to any privacy preference signal and any privacy flag, including those contained in consent cookies. The general idea is to compare a site's privacy flags in third party web requests before and after it received a user's privacy preference signal-enabled opt out request. Similarly to the US Privacy String of the IAB, Google provides sites that are implementing Google Ads the RDP flag [32]. Google explicitly explains that “partners who have implemented the Global Privacy Control may choose to enable restricted data processing [RDP] when they receive a GPC opt-out signal” [32]. Enabling this setting is the responsibility of the site operator. Once RDP is enabled, Google will limit its use of the supplied data [32]. The same is true when Meta's LDU flag

is enabled [52]. In such case, Meta will process data as a service provider and will limit use of the flagged data, in particular, as to tracking and measurement of ad effectiveness [52]. Figure 5 shows an illustrative overview of the general methodology.

As various CMPs implement GPC, their consent cookies may also contain privacy flags. For example, site operators that implement OneTrust [57] can set the US Privacy String via the OptanonConsent cookie. They can set the cookie value to `isGpcEnabled=1` to respect GPC or alternatively to `isGpcEnabled=0`. Such settings can serve as indicators for whether a site is respecting people's opt out choices. However, it should be noted that not all CMPs provide such compliance setting, and even if they do, not all site operators choose to use their CMPs for such purpose. In those cases setting a privacy flag to opt out in a CMP cookie after receiving a GPC signal would only indicate that the GPC signal was received but not that it was passed on to downstream third party sites. In essence, in those cases enabling a privacy flag in a CMP cookie upon receiving a GPC signal would only indicate that the CMP functionality for detecting a certain privacy preference signal was correctly implemented but not that the signal was passed on.

4.1.3 Limitations of Detecting GPC Non-compliance. There are various limitations for detecting GPC non-compliance. Some are based on the online ad ecosystem; others are inherent in our methodology. First, implementing functionality for detecting GPC signals is not sufficient for a site to be compliant. Rather, sites would also need to implement functionality for acting on the receipt of GPC signals, e.g., by changing the US Privacy String. Seeing string characters change from N to Y in the third position of the US Privacy String hints that a site has implemented such mechanisms, however, does not conclusively prove that a user is not being tracked. In many cases, the web requests to third parties will still include the data, and the recipients could simply disregard the opt out configuration of the passed on privacy flag. Ultimately, much of the current online ad ecosystem is built on the assumption that the recipients of data will act in compliance with the law. The GPC draft standard gives sites the option to indicate with `/.well-known/gpc.json` resource that they respect GPC [30]. However, the inclusion of such resource by sites respecting GPC is not mandatory. It is a promise rather than an enforceable mechanism.

Further, with our methodology, it is only possible to check the privacy flag passed from a first party to a third party but not between third parties, for example, when personal information is sold from one third party to another third party backend to backend. How ad networks and other third parties are sharing data at the backend remains non-transparent. Also, we rely on sites' Do Not Sell links and US Privacy String implementations to determine whether the CCPA is applicable to them. Thus, if a site does not have a Do Not Sell link despite being subject to the CCPA, additional effort would be required to determine the applicability of the CCPA [77]. Similarly, if a site does not implement the US Privacy String, our methodology would not be able to identify GPC non-compliance. Though, it should be noted that the US Privacy String is a commonly used privacy flag. As of the end of August 2022 the lead generation service BuiltWith identified more than three million sites on the web that have some indication in their code for potentially implementing the US Privacy String in form of

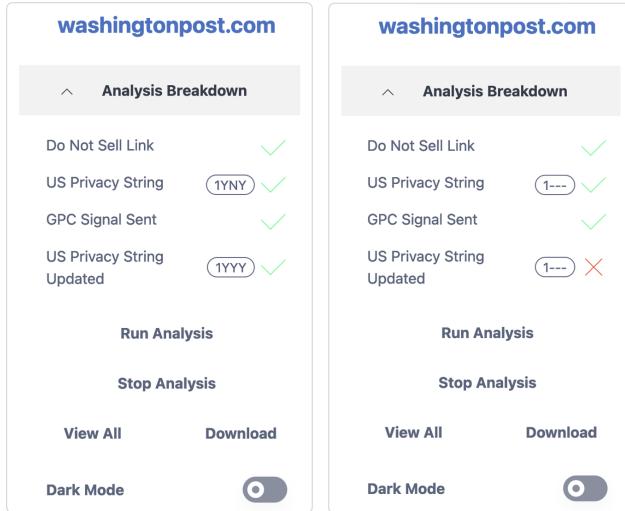


Figure 6: Labels from our browser extension showing analysis results for `washingtonpost.com` given a California IP address (left) and a non-California IP address (right). The site has a Do Not Sell link and for California its initial US Privacy String is 1NY. After receipt of a GPC signal, the site opts out people from California as indicated by the US Privacy String changing to 1YYY. Outside of California, the site does not apply the CCPA as indicated by 1- - -.

```

1 ((California|CA).?Resident).{0,10}((Do.?Not|Don.?t).?Sell)|  

2 (Do.?Not|Don.?t).?Sell.?{My}?.?(Personal)?.?(Information|Info|Data)

```

Figure 7: The JavaScript regular expression developed in the validation phase and used in our extension to identify Do Not Sell links.

the USPAPI [10].²⁴ As our methodology leverages artifacts of the current online ad ecosystem it presumes that site operators participate in it and implement its artifacts properly. Also, the US Privacy String is specific to the CCPA. Thus, sites may behave differently depending on whether they are accessed from inside or outside of California. For our analysis we used a VPN set to a California IP address (§5). Sites could have detected that we were not located in California, which could have impacted our results.

4.2 Implementation

We implemented a proof-of-concept browser extension for Firefox that performs the four steps of our compliance analysis methodology automatically (§4.1.1) and returns analysis results in label format as shown in Figure 6.²⁵

4.2.1 A Proof-of-concept Browser Extension. When the user starts an analysis, our extension searches for a Do Not Sell link on the visited website based on a regular expression match (Figure 7). We consider sites subject to the CCPA opt out requirements and, thus,

²⁴The operators of BuiltWith added functionality for detecting such sites after we contacted them with this suggestion. They also added functionality for detecting whether a US Privacy String value is returned as well as functionality for detecting sites potentially implementing Google's RDP or Meta's LDU privacy flags.

²⁵Our code is available at <https://github.com/privacy-tech-lab/gpc-optmeowt>.

subject to GPC if they have both a Do Not Sell link and US Privacy String. Neither would be necessary otherwise. Our extension identifies Do Not Sell links by intercepting the web requests the browser makes via the `webRequest` API. Dynamically identifying Do Not Sell links in web requests is less error-prone than doing so statically in program code as the latter may miss links injected at runtime via scripts or other resources. Our extension then calls the USPAPI, checks its return value, and checks all cookies in the intercepted web requests for the presence of a US Privacy String to capture its value, if any. Then, it sends a GPC signal, reloads the site, and checks again for the US Privacy String. If a site respects GPC, i.e., after receiving the GPC signal the third character of the US Privacy String is a Y, our extension indicates GPC compliance; otherwise, it indicates non-compliance (Figure 6). This four-pronged analysis can be generally applied to other privacy preference signals as well: by checking applicability of a law, testing a site's status, sending the signal, and retesting the site's status, it can be determined if a signal is respected or not.

4.2.2 Accuracy of Non-compliance Detection. After initially tuning our implementation on a validation set of 80 sites (Validation Set), we evaluated how accurately our extension works on a held-out test set of 100 sites (Test Set). We selected the sites for our Validation and Test Sets randomly from the top 1,000 sites on the Tranco list [60]. We excluded sites that would likely not be subject to the CCPA, i.e., sites with top level domains that are typically not used for US sites, such as .cn. We wanted to create datasets with some sites that have Do Not Sell links and US Privacy Strings and some sites that do not have such links and Strings. Such distributions are important because we leverage Do Not Sell links and US Privacy Strings as identifiers for whether a site is subject to the CCPA's opt out rules. We obtained ground truth data for the sites in our Validation and Test Sets by performing our analysis method manually. In detail, for each site we performed the following steps:

- (1) Visual inspection for a Do Not Sell link scrolling to the end of the site as necessary
- (2) Detection of the US Privacy String value by calling the USPAPI from the Firefox Web Console and checking the site's cookies via the Firefox Network Monitor
- (3) Sending a GPC signal to the site
- (4) Repeating the second step to detect changes in the US Privacy String, if any

We declared a US Privacy String present on a site if the USPAPI returned a US Privacy String value or we found such in a `us_privacy` or similarly named cookie. We did not consider a null return value from a USPAPI call as a valid US Privacy String value. As sites are only required to honor GPC for California residents we performed all our analyses while being connected to a VPN with a Los Angeles IP address.

Our extension ran successfully and correctly on 89/100 sites of the Test Set (Figure 8). 4/100 sites prevented our extension's inline JavaScript injection due to their Content Security Policy (CSP).²⁶ Excluding the sites not analyzable due to their CSP, our

²⁶After this finding we added functionality to our extension that disables Firefox's CSP check. While such setting is certainly not advisable for everyday browsing, it is tolerable in the context of our use case, i.e., GPC compliance analysis. Users of our extension will receive a respective notification.

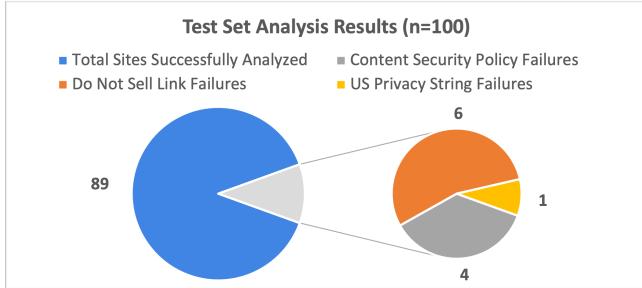


Figure 8: GPC compliance analysis results for our extension running on the 100 sites of our Test Set. The majority of analysis failures occurred due to incorrect Do Not Sell link identification (6/100) and sites' Content Security Policy blocking our extension (4/100). One site was incorrectly identified as having a US Privacy String while that was actually not the case per its ground truth (1/100).

Analysis Item	TP, FP, TN, FN	Tot	Acc	Prec	Rec	F-1
Do Not Sell Link Found	49, 4, 41, 2	96	0.96	0.96	0.92	0.94
USPS Found Before GPC Sent	31, 1, 64, 0	96	0.99	1	0.97	0.98
GPC Sent	96, 0, 0, 0	96	1	1	1	1
USPS Found after GPC Sent	31, 1, 64, 0	96	0.99	1	0.97	0.98
USPS Opt Out after GPC Sent	4, 0, 27, 0	31	1	1	1	1
USPS Change to Opt Out after GPC Sent	3, 0, 0, 0	3	1	1	1	1

Table 1: Detailed GPC compliance analysis results of our extension for a subset of 96/100 sites of our Test Set. The subset excludes 4/100 sites for which the analysis failed due to their Content Security Policy. For all 31/96 sites with US Privacy Strings, the existence of the Strings and correct values were detected both before and after sending the GPC signal. (USPS = US Privacy String, TP = True Positives, FP = False Positives, TN = True Negatives, FN = False Negatives, Tot = Total, Acc = Accuracy, Prec = Precision, Rec = Recall, F-1 = F-1 score.)

extension's identification of Do Not Sell links on the remaining subset of 96/100 sites reached an F-1 score of 0.94 when compared against the manual ground truth analysis results (Table 1). This performance indicates the overall suitability of dynamic web request analysis in combination with regular expressions for Do Not Sell link identification. False positives can occur due to unrelated text on sites mentioning "Do Not Sell" and false negatives due to the use of Do Not Sell images instead of character strings. Natural language processing and machine learning techniques could be fruitful techniques for tackling the former challenge and optical character recognition for the latter.

Correctly identifying Do Not Sell links is not a trivial task [77]. Indeed, it turns out to be the largest error source of our extension's analysis performance. On the other hand, identifying US Privacy Strings and their values — whether before or after sending the GPC signal — is much less error-prone. In our subset of 96/100 Test Set sites, 31/96 sites implement the US Privacy String. With one false positive and an F-1 score of 0.98 our extension reliably identified the US Privacy String both in terms of its existence on a site as well as its value. 3/31 of the sites with a US Privacy String changed its third character to a Y after receiving a GPC signal and

1/31 opted out users already prior to receiving GPC signals, all of which were correctly analyzed by our extension. We also note that correctly identifying the US Privacy String in the USPAPI is more important than in cookies. For the 31 Test Set sites implementing the US Privacy String 15/31 use the USPAPI, 2/31 use cookies, and 14/31 use both the USPAPI and cookies.

5 GPC COMPLIANCE OF TOP WEBSITES

Our results suggest that people would benefit greatly from the availability of GPC in browsers and browser extensions to efficiently and effectively exercise their opt out rights per the CCPA and other new privacy laws (§3). However, people's opt out choices would not mean much meaning if sites do not respect GPC at a broad scale. Thus, we perform a GPC compliance enforcement analysis (§5.1) and show that our GPC compliance analysis can be scaled by using our browser extension in combination with a web crawler (§5.2). Scaling GPC compliance analysis is important to make websites' privacy practices transparent and broadly enforce GPC.

5.1 Enforcing GPC Compliance

For our GPC compliance enforcement analysis, we manually analyzed GPC applicability and compliance for the 180 websites in our Validation and Test Sets combined (Combined Set). The GPC compliance analysis results are based on the ground truth for each of the 180 sites. All sites in the Combined Set were accessed via a VPN with a Los Angeles IP address to trigger the applicability of the CCPA opt out requirements for sites subject to it.

5.1.1 Sites' Do Not Sell Links and US Privacy Strings. 121/180 of the sites in the Combined Set have a Do Not Sell link (Figure 9A). 59/121 sites with Do Not Sell links also have US Privacy Strings. For the remaining 62/121 sites their operators may not be IAB members or they may have decided to not implement the US Privacy String. Interestingly, not all sites that implement the US Privacy String also implement a Do Not Sell link. 5/59 sites do not implement a Do Not Sell link but do implement the US Privacy String. These implementations are inconsistent: if sites are subject to the CCPA, they should have a Do Not Sell link in addition to the US Privacy String; if they are not subject to the CCPA, they should not implement either. This inconsistency indicates that some site operators struggle with CCPA compliance. Their inconsistent implementations may leave people confused and unable to opt out.

The majority of the 64 sites that implement the US Privacy String do so either via the USPAPI or via both the USPAPI and cookies (Figure 9B). However, the implementation with just cookies is infrequent and only occurs on 3/64 sites. It could be that site operators do not want to rely on the US Privacy String cookie implementation as browser vendors will eventually phase out third-party cookies [63]. Another reason in favor of implementing the US Privacy String via the USPAPI is that the GPC choice of a user will be available quicker than via cookies, which, in turn, enables a site to make quicker decisions as to which ads to serve, if any.

5.1.2 A GPC Compliance Analysis Case Study. Whichever technology sites select for implementing the US Privacy String, only few respect GPC. Initially, only 8/64 sites changed the third position of the US Privacy String to a Y and opted us out (Figure 9C). However,

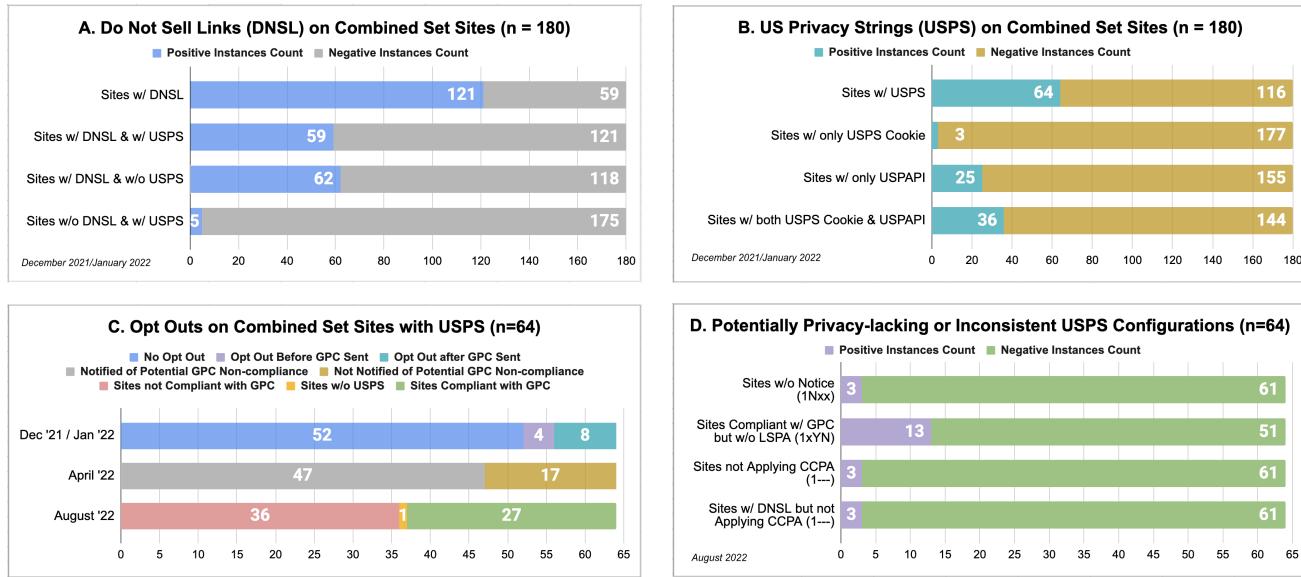


Figure 9: Ground truth statistics covering Do Not Sell links, US Privacy Strings, and GPC compliance for the 180 sites in our Combined Set.

an additional 4/64 sites also had us already opted out before we sent a GPC signal. 52/64 sites did not react to the GPC signal we sent and did not set the US Privacy String to a Y in the third position. The 52 sites include two sites that opted us out before sending a GPC signal, however, after receiving the GPC signal opted us back in. This behavior may have been based on an incorrect implementation of the US Privacy String. We contacted the two site operators but did not hear back. In our subsequent analysis, we found that one site no longer had the US Privacy String implemented and the other no longer opted us out – neither before nor after receiving a GPC signal. Both sites are included in the set of 47/52 sites that we notified via email in April 2022 about their potential GPC non-compliance that we detected in December 2021 and January 2022. For 5/52 sites that were potentially non-compliant we could not find email contact information and, thus, did not contact them. We made clear that our email may serve as a notice for triggering the 30-day cure period per the CCPA to prepare potential enforcement actions by the OAG.²⁷ In June 2022 we sent a follow-up email with additional GPC implementation guidance to site operators from whom we had not heard back about their GPC implementation plans. We again offered our help in implementing GPC.

In all our communications with the site operators we truthfully disclosed our identity, our affiliation, our findings about their site, and the purpose for why we are contacting them. All emails were sent from the institutional email account of the first author and signed by the first author. In all emails we offered site operators our help in implementing GPC or ask any questions they may have. For all 47 sites whose operators we contacted we had manual ground truth analysis results indicating their non-compliance. 22/47 site operators responded substantively to our inquiries:

- 9/22 confirmed that they have implemented GPC

- 6/22 replied that they are considering implementing GPC
- 4/22 replied that they rely on their CMP for GPC compliance
- 1/22 clarified that GPC does not apply to their site
- 1/22 referred us to their Do Not Sell link
- 1/22 limited their response to the US Privacy String

All responses to our inquiries were friendly in tone except for one in which the operator accused us of a “poorly veiled attempt at blackmail” and asked us rhetorically whether we would be content if they share our email with our “higher ups.” We did not engage in further communication with this operator. We were also contacted by attorneys of a large US law firm who explained that several of their clients had received our emails. They did not disclose who their clients were. We scheduled an informational call with them that took place in an agreeable atmosphere. We do not know what advice they gave to their clients based on the explanations we provided in our call.

In August 2022, we re-tested all non-compliant sites and found that 15/47 had become compliant in the meantime. 14/15 sites implemented GPC and, as described, one site operator had clarified that they are not selling personal information and, thus, are not subject to the CCPA’s Do Not Sell provisions. As of August 2022, a total of 27/64 sites that implement the US Privacy String were compliant and respected GPC.²⁸ The majority of the sites that respect GPC are national and global information, news, and media sites. It appears that larger publishers are more likely to respect GPC than smaller ones. Possibly, the stronger first party relationship they may have with their users compared to smaller sites makes it easier for them to adopt privacy-preserving ad business models. The availability of more resources to react to new legal developments and the higher risk of being in the news for non-compliant privacy practices may play a role as well.

²⁷ California Civil Code §1798.150(b).

²⁸ Appendix A.3 lists the 27 compliant sites.

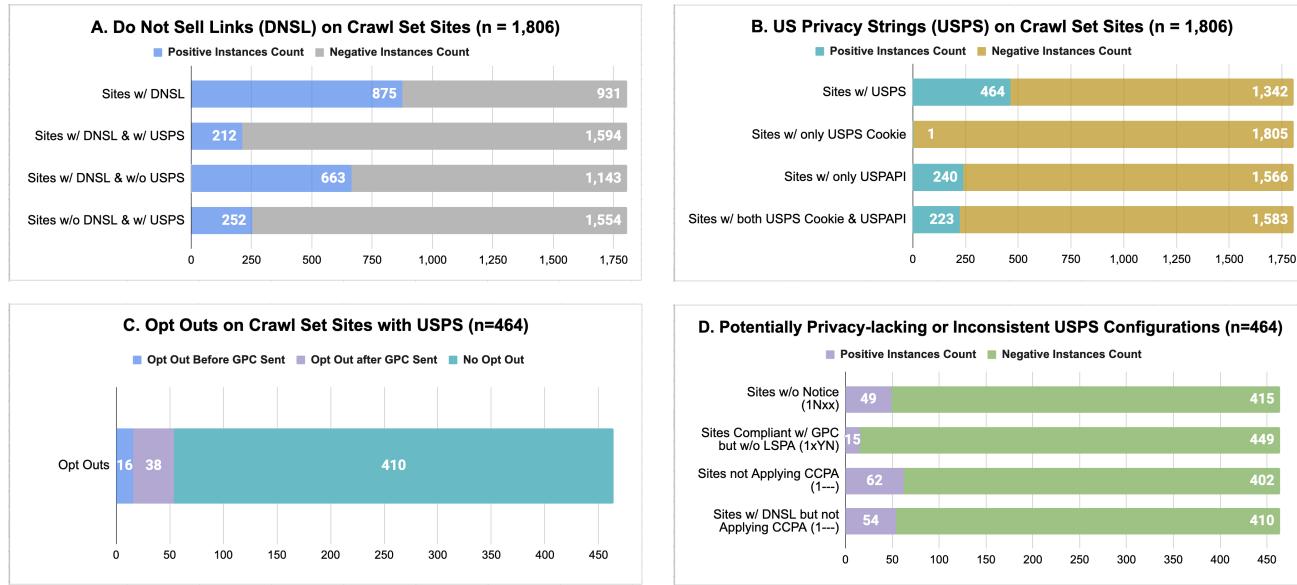


Figure 10: Statistics covering Do Not Sell links, US Privacy Strings, and GPC compliance for the 1,806 sites in our Crawl Set (all August 2022).

5.1.3 Reporting Our Findings. For non-compliant sites as of August 2022 our findings provide evidence for its operators breaking the law. Such non-compliance may not always be intentional but can be due to a lack of knowledge or an implementation mistake, among others. Per the CCPA and CCPA Regulations, respecting opt outs via GPC is mandatory for sites to which these laws apply [73]. To that effect, the OAG recently began enforcing GPC against non-compliant businesses [71]. Notably, the OAG brought an enforcement action and entered into a settlement agreement with fashion retailer Sephora for failing to disclose that it was selling personal information and not processing opt outs via “global privacy controls” [72]. Such enforcement actions are critical for the broad adoption of GPC, which we hope to advance with our work.

We saw it as our responsible disclosure obligation to first notify site operators about our findings and give them the opportunity and time to implement working GPC functionality on their sites.²⁹ We offered them our help in doing so and asked them to get in touch if they had any questions or were interested in further discussing this matter with us. We emphasized our willingness to help but also made clear that we would notify the OAG with the findings of our research. We have not disclosed any individual finding of non-compliance publicly and will not do so in the future. One reason is that such publication could interfere with potential GPC enforcement actions by the OAG, to which we are not privy. To ensure that all applicable legal and ethical considerations of our report to the OAG were considered we conferred internally with our institution’s IRB staff, legal counsel, communications staff, university administrators, and faculty colleagues about the procedure and potential ramifications.

5.1.4 US Privacy String Values in Detail. The values of the US Privacy Strings exposed by the websites in our Combined Set provide

further insight into how opt outs are being processed. 3/64 sites did not provide notice – neither before nor after receiving our GPC signals – as indicated by their 1Nxx US Privacy String (Figure 9D).³⁰ Also, 13/64 sites configured the US Privacy String as 1xYN. This configuration indicates that they respect GPC but are not operating under the IAB’s Limited Service Provider Agreement (LSPA) [37]. If they would operate under the LSPA, they would only be allowed to transact with signatories of the LSPA and any of their subproviders. However, as they designated the transaction outside the scope of the LSPA, they are not limited in such a way. Thus, users’ opt outs could have only limited effect. Further, 3/64 sites, all of which are non-compliant, indicated that the CCPA is not applicable. They configured their US Privacy String as 1--. However, the same 3/64 sites also displayed a Do Not Sell link potentially contradicting their US Privacy String configuration. Especially, the third character position cannot be unknown, i.e., must never include a hyphen, if the CCPA applies [40]. These are further instances of inconsistent CCPA implementations with the potential of confusing people or not honoring their opt out choices as required by law.

5.2 Scaling GPC Compliance Analysis

We developed a web crawler for our extension and analyzed 1,806 websites for GPC applicability and compliance (Crawl Set).

5.2.1 Methodology. We implemented our crawler to run on Puppeteer for Firefox Nightly [33].³¹ The analysis results for the Crawl Set are provided here as they were returned by the automated analysis from our browser extension without ground truth verification. We selected the sites for the Crawl Set based on a set of sites that had some indication in their code for implementing the USPAPI

³⁰Notice is required per California Civil Code §1798.115(d).

³¹Our code is available at <https://github.com/privacy-tech-lab/gpc-optmeowt>.

²⁹The responsibility for the reporting lies with the first author.

according to the lead generation service BuiltWith [10]. Specifically, we selected the first 2,000 sites from the set according to their Tranco list rank. We excluded sites contained in our Combined Set. To obtain an unbiased assessment of the performance of our crawler we also excluded sites from the set of 2,000 sites that we had selected randomly for tuning our crawler during its development. After these exclusions, the Crawl Set contained 1,861 sites.

Our crawler ran successfully on 1,806/1,861 (97%) sites. We implemented it in non-headless mode leveraging Puppeteer's Keyboard API for starting and stopping the analysis of a site via a keyboard command on a virtual keyboard. Pressing the virtual keyboard can lead to errors on sites with text boxes as the keyboard commands are entered into the text boxes instead of being executed. We minimized occurrences of this error by letting our crawler issue `shift+tab` commands first. A few sites also failed to load independently of our crawler. We crawled the set of 1,861 sites on a Mac Mini set to a California IP address via a VPN in three batches, each of which ran without interruption or crash. Crawling the set of 1,861 sites took a total of 21 hours, 22 minutes, and 7 seconds with a mean analysis time of 41.3 seconds per site. The analysis of each site includes two 15-second timeouts for waiting for the site to load before querying the US Privacy String before and after sending the GPC signal. In addition, 1-second timeouts are set to wait for each keyboard command, and a 5-second timeout is set before closing the analyzed site and starting the analysis of the next site.

5.2.2 Only Few Sites Respect GPC. Given the accuracy of our extension's GPC analysis (Table 1), the major trends of GPC applicability and compliance of sites in the Crawl Set as a whole can be assumed to be represented faithfully while analysis errors for individual sites cannot be ruled out. 464 sites had a US Privacy String (Figure 10B). Of those, only 212 had a Do Not Sell link (Figure 10A). Thus, 252/464 sites are potentially non-compliant with the CCPA as they inherently declare through their US Privacy String that they are selling data but do not offer a Do Not Sell link opt out option. This proportion of sites is much larger than in the Combined Set where only 5/64 sites with a US Privacy String did not have a Do Not Sell link (Figure 9A). This result shows the potential of using the US Privacy String for identifying compliance violations at scale. Further, with only 54/464 (12%) of the sites with US Privacy String opting us out — 16/54 sites before sending a GPC signal and 38/54 after — GPC non-compliance appears widespread (Figure 10C). Also, for three sites, which are not counted in the set of 54 compliant sites, we were opted back in after initially having been opted out, just as for the two sites in the Combined Set. Such behavior is worrisome as it may give people a wrong sense of privacy.

5.2.3 US Privacy String Values in Detail. Similarly as for the Combined Set, for a subset of sites in the Crawl Set the values of the US Privacy String are dubious (Figure 10D). 49/464 sites did not provide notice before or after receiving our GPC signals as indicated by their 1Nxx US Privacy String, 15/464 sites configured their US Privacy String as 1xYN indicating that they are not operating under the LSPA [37] and, thus, with possible lesser privacy protections, and 54/464 sites indicated that the CCPA is not applicable by configuring their US Privacy String as 1- - - while at the same time contradicting their statement by displaying a Do Not Sell link.

6 CONCLUSIONS

GPC holds the promise of broadly empowering people to efficiently and effectively opt out of web tracking and convey their choices accurately and intentionally to the sites they visit. To achieve this goal we need further progress in usability and enforceability.

6.1 Usability

While the law is catching up with the dynamic development of the web, privacy rights will not matter much if they remain too impractical to exercise for anyone but the most dedicated privacy enthusiasts. Our usability survey results suggest that people understand what GPC does and that they would make use of GPC if it were available to them. To empower people to exercise their privacy rights, it is critical to implement GPC in form of usable software. Notably, as 89% of participants in our survey would send GPC signals to all or most sites they visit, any GPC implementation should provide for a universal GPC setting. If people are made aware of such setting and how to change it, for example, during the setup of a new browser, GPC should be turned on by default to further enhance usability.

6.2 Enforceability

Enforceability starts with transparency. Generally, people have no way of knowing whether their opt outs or other privacy choices are respected by the sites they visit. Currently, non-compliant sites can largely remain undetected without much risk of being exposed. Thus, it is important to make a site's handling of GPC and other privacy preference signals visible. Transparently disclosing a site's behavior serves as a strong motivator for compliance. We described a methodology for surfacing whether a site propagates opt out choices to ad networks and other downstream providers it integrates as required per the CCPA after receiving a GPC signal. This approach is generalizable and can also be applied to other privacy flags and privacy preference signals sent per the GDPR, for example.

It is critical that GPC is mandatory. GPC will not have a broad impact if sites are not required to respect it. This lesson was learned from DNT and is being applied in the CCPA Regulations with the OAG enforcing GPC for California consumers [71–73]. Regulators in Colorado and Connecticut are also on their way to make GPC mandatory per their state privacy laws. If sites are required to respect GPC, browser vendors would have stronger incentives to implement GPC functionality. Otherwise, they may worry that GPC will just increase a browser's fingerprinting surface without much privacy gain possibly even resulting in a privacy net negative.

In the long term, the privacy challenges of the web ecosystem are best addressed by integrating privacy protections directly into the architecture of the web. This is the reason why it is so important to pursue standardization of GPC and other privacy preferences signals at the W3C and other standards bodies. To the extent that privacy preference signals become part of the web, they will become automatically enforced due to their nature as standards. As the CCPA Regulations show, regulators can play an important role in this process by delegating the technical workings of privacy rights in their regulations to privacy standards that define the technical means for executing the law.

ACKNOWLEDGMENTS

We would like to thank our shepherd Tobias Urban and the anonymous reviewers for their valuable feedback. We appreciate their extraordinary effort and diligence in reviewing our paper. Wesleyan student Kate Hausladen did an excellent job helping with the software development, for which we thank her very much. We are grateful to the National Science Foundation (Award #2055196) and to the Alfred P. Sloan Foundation (Program in Universal Access to Knowledge) for their support of this research. We also thank Wesleyan University, its Department of Mathematics and Computer Science, and the Anil Fernando Endowment for their additional support. Andrew Rogers and Gary Brewer of BuiltWith added privacy flag identification functionality to their service and provided us with a free Pro account, for which we are grateful. Conclusions reached or positions taken are our own and not necessarily those of our financial supporters, its trustees, officers, or staff.

AVAILABILITY OF ARTIFACTS

Our code is available at <https://github.com/privacy-tech-lab/gpc-optmeowt>.

REFERENCES

- [1] Advanced Data Protection Control (ADPC). 2022. <https://www.dataprotectioncontrol.org/>. Accessed: March 7, 2023.
- [2] Rakesh Agrawal, Jerry Kiernan, Ramakrishnan Srikant, and Yirong Xu. 2003. An XPath-based preference language for P3P. In *WWW*. ACM, New York, NY, USA, 629–639. <https://doi.org/10.1145/775152.775241>
- [3] Apple Developer. 2022. App Tracking Transparency. <https://developer.apple.com/documentation/apptrackingtransparency>. Accessed: March 7, 2023.
- [4] Rebecca Balebako, Pedro G. Leon, Richard Shay, Blase Ur, Yang Wang, and Lorrie Faith Cranor. 2012. Measuring the effectiveness of privacy tools for limiting behavioral advertising. https://www.researchgate.net/publication/267705080_Measuring_the_Effectiveness_of_Privacy_Tools_for_Limiting_Behavioral_Advertising. In *Web 2.0 Workshop on Security and Privacy*. IEEE, San Francisco, CA, USA, 1–10.
- [5] Vinayak Bannihatti Kumar, Roger Iyengar, Namita Nisal, Yuanyuan Feng, Hana Habib, Peter Story, Sushain Cherivirala, Margaret Hagan, Lorrie Cranor, Shomir Wilson, Florian Schaub, and Norman Sadeh. 2020. Finding a Choice in a Haystack: Automatic Extraction of Opt-Out Statements from Privacy Policy Text. In *WWW*. ACM, New York, NY, USA, 1943–1954. <https://doi.org/10.1145/3366423.3380262>
- [6] Xavier Becerra. 2021. <https://twitter.com/AGBecerra/status/1354850758236102656>. accessed: March 7, 2023.
- [7] Belgian Data Protection Authority. 2022. Decision on the merits 21/2022 of 2 February 2022. <https://www.autoriteprotectiondonnees.be/publications/decision-quant-au-fond-n-21-2022-english.pdf>. Accessed: March 7, 2023.
- [8] Robin Berjon. 2021. GPC under the GDPR. <https://berjon.com/gpc-under-the-gdpr>. Accessed: March 7, 2023.
- [9] Sophie C. Boerman, Sanne Kruikemeier, and Frederik J. Zuiderveen Borgesius. 2021. Exploring Motivations for Online Privacy Protection Behavior: Insights From Panel Data. <https://journals.sagepub.com/doi/10.1177/0093650218800915>. *Communication Research* 48, 7 (2021), 953–977.
- [10] BuiltWith. 2022. US Privacy User Signal Mechanism Usage Statistics. <https://trends.builtwith.com/widgets/US-Privacy-User-Signal-Mechanism>. Accessed: March 7, 2023.
- [11] Matt Burgess. 2022. Google Has a New Plan to Kill Cookies. People Are Still Mad. <https://www.wired.com/story/google-floc-cookies-chrome-topics/>. Accessed: March 7, 2023.
- [12] Simon Byers, Lorrie Faith Cranor, Dave Kormann, and Patrick McDaniel. 2005. Searching for privacy: design and implementation of a P3P-enabled search engine. In *PETS*. Springer, Berlin, Heidelberg, Germany, 314–328. https://doi.org/10.1007/11423409_20
- [13] California Legislative Information. 2013. California Online Privacy Protection Act. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=8.&chapter=22.&lawCode=BPC. Accessed: March 7, 2023.
- [14] California State Legislature. 2018. Assembly Bill No. 375. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5. Accessed: March 7, 2023.
- [15] Anthony Chavez. 2022. Introducing the Privacy Sandbox on Android. <https://blog.google/products/android/introducing-privacy-sandbox-android/>. Accessed: March 7, 2023.
- [16] Consumer Reports. 2022. Data Rights Protocol (DRP). <https://github.com/consumer-reports-digital-lab/data-rights-protocol>. Accessed: March 7, 2023.
- [17] Lorrie Cranor. 2018. P3P Compact Policy Cross-Reference. https://web.archive.org/web/20140908033337/http://compactprivacypolicy.org/compact_token_reference.htm. Accessed: March 7, 2023.
- [18] Lorrie Faith Cranor, Brooks Dobbs, Serge Egelman, Giles Hogben, Jack Humphrey, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, Joseph M. Reagle, Matthias Schunter, David A. Stampley, and Rigo Wenning. 2006. The Platform for Privacy Preferences 1.1 (P3P1.1) Specification. <https://www.w3.org/TR/P3P11/>
- [19] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006. User interfaces for privacy agents. *ACM Trans. Comput.-Hum. Interact.* 13, 2 (June 2006), 135–178. <https://doi.org/10.1145/1165734.1165735>
- [20] Lorrie Faith Cranor, Marc Langheinrich, and Massimo Marchiori. 2002. A P3P Preference Exchange Language 1.0 (APPEL 1.0). <http://www.w3.org/TR/2002/WD-P3P-preferences-20020415>.
- [21] Lorrie Faith Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph M. Reagle. 2002. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. <https://www.w3.org/TR/P3P1>.
- [22] DAA. 2020. DAA CCPA Opt Out Tool for the Web. https://digitaladvertisingalliance.org/DAA_style/ADS/CCPA_Opt_Out_Tool_Technical_Description.pdf. Accessed: March 7, 2023.
- [23] DAA. 2022. YourAdChoices. <https://youradchoices.com/>. Accessed: March 7, 2023.
- [24] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN>. Accessed: March 7, 2023.
- [25] Facebook. 2020. Data Processing Options for Users in California. <https://developers.facebook.com/docs/marketing-apis/data-processing-options>. Accessed: March 7, 2023.
- [26] Glenn Fleishman. 2019. How the tragic death of Do Not Track ruined the web for everyone. <https://www.fastcompany.com/90308068/how-the-tragic-death-of-do-not-track-ruined-the-web-for-everyone>. Accessed: March 7, 2023.
- [27] General Assembly of the State of Colorado. 2021. Colorado Privacy Act. https://leg.colorado.gov/sites/default/files/documents/2021A/bills/2021a_190_rer.pdf. Accessed: March 7, 2023.
- [28] General Assembly of Virginia. 2021. Virginia Consumer Data Protection Act. <https://lis.virginia.gov/cgi-bin/legp604.exe?212+sum+HB2307>. Accessed: March 7, 2023.
- [29] Global Privacy Control (GPC) Group. 2022. <https://globalprivacycontrol.org/>. Accessed: March 7, 2023.
- [30] Global Privacy Control (GPC) Group. 2022. Global Privacy Control (GPC). <https://globalprivacycontrol.github.io/gpc-spec/>. Accessed: March 7, 2023.
- [31] Global Privacy Control (GPC) Group. 2022. Interacting with Global Privacy Control. <https://global-privacy-control.glitch.me/>. Accessed: March 7, 2023.
- [32] Google. 2022. Helping advertisers comply with CCPA in Google Ads. <https://support.google.com/google-ads/answer/9614122?hl=en>. Accessed: March 7, 2023.
- [33] Google. 2022. Puppeteer. <https://pptr.dev/>. Accessed: March 7, 2023.
- [34] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. 2021. Conflicting Privacy Preference Signals in the Wild. <https://arxiv.org/pdf/2109.14286.pdf>. Accessed: March 7, 2023.
- [35] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. 2021. Privacy Preference Signals: Past, Present and Future. In *PETS (2021)*. Sciendo, Innsbruck, Austria, 249–269. <https://petsymposium.org/2021/files/papers/issue4/pets-2021-0069.pdf>
- [36] Soheil Human, Harshvardhan J. Pandit, Victor Morel, Cristiana Santos, Martin Degeling, Arianna Rossi, Wilhelmina Botes, Vitor Jesus, and Irene Kamara. 2022. Data Protection and Consenting Communication Mechanisms: Current Open Proposals and Challenges. <https://harshp.com/research/publications/051-Consenting-Communication-Mechanisms>. Accessed: March 7, 2023.
- [37] IAB. 2019. Limited Service Provider Agreement. <https://web.archive.org/web/20220817164750/https://www.iabprivacy.com/lspa-2019-12.pdf>. Accessed: March 7, 2023.
- [38] IAB. 2020. IAB CCPA Compliance Framework For Publishers & Technology Companies. <https://iabtechlab.com/standards/ccpa/>. Accessed: March 7, 2023.
- [39] IAB. 2020. US Privacy User Signal Mechanism “USP API”. <https://github.com/InteractiveAdvertisingBureau/USPrivacy/blob/master/CCPA/USP%20API.md>. Accessed: March 7, 2023.
- [40] IAB. 2021. US Privacy String. <https://github.com/InteractiveAdvertisingBureau/USPrivacy/blob/master/CCPA/US%20Privacy%20String.md>. Accessed: March 7, 2023.
- [41] IAB. 2022. Propose a GPC extension to OpenRTB. <https://github.com/InteractiveAdvertisingBureau/openrtb/pull/99>. Accessed: March 7, 2023.
- [42] IAB Europe. 2021. IAB TCF - Transparency & Consent Framework. <https://iabeurope.eu/transparency-consent-framework/>. Accessed: March 7, 2023.

[43] IAB Europe. 2022. Transparency and Consent String with Global Vendor & CMP List Formats. <https://github.com/InteractiveAdvertisingBureau/GDPR-Transparency-and-Consent-Framework/blob/master/TCFv2/IAB%20Tech%20Lab%20-%20Consent%20string%20and%20vendor%20list%20formats%20v2.md>. Accessed: March 7, 2023.

[44] IAB Tech Lab. 2022. Global Privacy Platform. <https://iabtechlab.com/gpp/>. Accessed: March 7, 2023.

[45] Kate Kaye. 2021. California Attorney General says popular, digital ad opt-outs from trade groups don't comply with CCPA. <https://digiday.com/media/california-attorney-general-says-popular-digital-ad-opt-outs-from-trade-groups-dont-comply-with-ccpa/>. Accessed: March 7, 2023.

[46] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "nutrition label" for privacy. In *SOUPS*. ACM, New York, NY, USA, Article 4, 12 pages. <https://doi.org/10.1145/1572532.1572538>

[47] Konrad Kollning, Anastasia Shuba, Max Van Kleek, Reuben Binns, and Nigel Shadbolt. 2022. Goodbye Tracking? Impact of iOS App Tracking Transparency and Privacy Labels. In *FAccT '22*. ACM, New York, NY, USA, 508–520. <https://doi.org/10.1145/3531146.3533116>

[48] Pedro Giovanni Leon, Lorrie Faith Cranor, Aleecia M. McDonald, and Robert McGuire. 2010. Token Attempt: The Misrepresentation of Website Privacy Policies Through the Misuse of P3P Compact Policy Tokens. In *WPES*. ACM, New York, NY, USA, 93–104. <https://doi.org/10.1145/1866919.1866932>

[49] Adam Lerner, Anna Kornfeld Simpson, Tadayoshi Kohno, and Franziska Roesner. 2016. Internet Jones and the Raiders of the Lost Trackers: An Archaeological Study of Web Tracking from 1996 to 2016. In *USENIX Security*. USENIX Association, Austin, TX, 997–1013. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/lerner>

[50] Maureen Mahoney. 2020. California Consumer Privacy Act: Are Consumer's Digital rights protected? https://advocacy.consumerreports.org/wp-content/uploads/2020/09/CR_CCPA-Are-Consumers-Digital-Rights-Protected_092020_vf.pdf. Accessed: March 7, 2023.

[51] Aleecia McDonald and Jon M. Peña. 2011. Track Gap: Policy Implications of User Expectations for the "Do Not Track" Internet Privacy Feature. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1993133. In *39th Research Conference on Communication, Information and Internet Policy*. TRPC, Arlington, VA, USA, 1–36.

[52] Meta. 2021. Limiting How Data is Used for People in California. <https://www.facebook.com/business/help/1151133471911882>. Accessed: March 7, 2023.

[53] Mozilla. 2021. Implementing Global Privacy Control. <https://blog.mozilla.org/netpolicy/2021/10/28/implementing-global-privacy-control/>. Accessed: March 7, 2023.

[54] Nevada State Assembly. 2019. Senate Bill No. 220. <https://www.leg.state.nv.us/NRS/NRS-603A.html>. Accessed: March 7, 2023.

[55] Rishab Nithyanand, Sheharbano Khattak, Mobin Javed, Narseo Vallina-Rodriguez, Marjan Falaharostegar, Julia Powles, Emiliano De Cristofaro, Hamed Haddadi, and Steven Murdoch. 2016. Ad-Blocking and Counter Blocking: A Slice of the Arms Race. <https://www.usenix.org/system/files/conference/foci16/foci16-paper-nithyanand.pdf>. In *FOCI*. USENIX Association, Ithaca, NY, USA, 1–7.

[56] Sean O'Connor, Ryan Nurwono, Aden Siebel, and Eleanor Birrell. 2021. (Un)Clear and (In)Conspicuous: The Right to Opt-out of Sale under CCPA. In *WPES*. ACM, New York, NY, USA, 59–72. <https://doi.org/10.1145/3463676.3485598>

[57] OneTrust. 2022. <https://www.onetrust.com/>. Accessed: March 7, 2023.

[58] Harshvardhan J. Pandit. 2021. GPC + GDPR: will it work? <https://harshp.com/research/blog/gpc-gdpr-can-it-work>. Accessed: March 7, 2023.

[59] David Pierce. 2021. DuckDuckGo's surprisingly simple plan to make the internet more private. <https://www.protocol.com/duckduckgo-ceo-interview>. Accessed: March 7, 2023.

[60] Victor Le Pochat, Tom van Goethem, Samaneh Tajalizadehkhooob, Maciej Korczynski, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. https://www.ndss-symposium.org/wp-content/uploads/2019/02/ndss2019_01B-3_LePochat_paper.pdf. In *NDSS*. Internet Society, VA, USA, 1–15.

[61] Prolific. 2022. Quickly find research participants you can trust. <https://www.prolific.co/>. Accessed: March 7, 2023.

[62] Robert W. Reeder, Patrick Gage Kelley, Aleecia M. McDonald, and Lorrie Faith Cranor. 2008. A user study of the Expandable Grid applied to P3P privacy policy visualization. In *WPES*. ACM, New York, NY, USA, 45–54. <https://doi.org/10.1145/1456403.1456413>

[63] Lawler Richard. 2022. Google delays blocking third-party cookies again, now targeting late 2024. <https://www.theverge.com/2022/7/27/23280905/google-chrome-cookies-privacy-sandbox-advertising>. accessed: March 7, 2023.

[64] Kanthashree Mysore Sathyendra, Shomir Wilson, Florian Schaub, Sebastian Zimmeck, and Norman Sadeh. 2017. Identifying the Provision of Choices in Privacy Policy Text. <https://aclanthology.org/D17-1294.pdf>. In *EMNLP*. ACL, Copenhagen, Denmark, 2764–2769.

[65] Ari Schwartz. 2009. Looking Back at P3P: Lessons for the Future. https://www.cdt.org/files/pdfs/P3P_Retro_Final_0.pdf. Accessed: March 7, 2023.

[66] Daniel Smullen, Yaxing Yao, Yuanyuan Feng, Norman Sadeh, Arthur Edelstein, and Rebecca Weiss. 2021. Managing Potentially Intrusive Practices in the Browser: A User-Centered Perspective. In *PETS (2021)*. Sciendo, Philadelphia, PA, USA, 500–527. <https://doi.org/10.2478/popets-2021-0082>

[67] Peter Snyder. 2020. Global Privacy Control, a new Privacy Standard Proposal. <https://brave.com/web-standards-at-brave/4-global-privacy-control/>. Accessed: March 7, 2023.

[68] Daniel J. Solove. 2021. The Myth of the Privacy Paradox. https://scholarship.law.gwu.edu/cgi/viewcontent.cgi?article=2738&context=faculty_publications. *George Washington Law Review* 89, 1 (2021), 1–51.

[69] State of California Department of Justice. 2020. California Consumer Privacy Act (CCPA) Final Statement of Reasons, Appendix E. <https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-fsor-appendix-e.pdf>. accessed: March 7, 2023.

[70] State of California Department of Justice. 2020. California Consumer Privacy Act Regulations. <https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/oal-sub-final-text-of-regs.pdf?>. Accessed: March 7, 2023.

[71] State of California Department of Justice. 2021. CCPA Enforcement Case Examples. <https://oag.ca.gov/privacy/ccpa/enforcement>. Accessed: March 7, 2023.

[72] State of California Department of Justice. 2022. Attorney General Bonta Announces Settlement with Sephora as Part of Ongoing Enforcement of California Consumer Privacy Act. <https://oag.ca.gov/news/press-releases/attorney-general-bonta-announces-settlement-sephora-part-ongoing-enforcement>. Accessed: March 7, 2023.

[73] State of California Department of Justice. 2022. California Consumer Privacy Act (CCPA) Frequently Asked Questions (FAQs). <https://oag.ca.gov/privacy/ccpa>. Accessed: March 7, 2023.

[74] State of Connecticut General Assembly. 2022. Connecticut Data Privacy Act. <https://www.cga.ct.gov/2022/AMD/S/pdf/2022SB-00006-R00SA-AMD.pdf>. Accessed: March 7, 2023.

[75] Utah State Legislature. 2022. Utah Consumer Privacy Act. <https://le.utah.gov/-/2022/bills/static/SB0227.html>. Accessed: March 7, 2023.

[76] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)Informed Consent: Studying GDPR Consent Notices in the Field. In *CCS*. ACM, New York, NY, USA, 973–990. <https://doi.org/10.1145/3319535.3354212>

[77] Maggie Van Nortwick and Christo Wilson. 2022. Setting the Bar Low: Are Websites Complying With the Minimum Requirements of the CCPA?. <https://petsympsymposium.org/2022/files/papers/issue1/popeps-2022-0030.pdf>. In *PETS*. Sciendo, Boston, MA, USA, 608–628.

[78] W3C. 2019. Tracking Preference Expression (DNT). <https://www.w3.org/TR/tracking-dnt/>. Accessed: March 7, 2023.

[79] W3C. 2019. Tracking Protection Working Group. <https://www.w3.org/2011/tracking-protection/>. Accessed: March 7, 2023.

[80] W3C. 2022. Consent Community Group. <https://www.w3.org/community/consent/>. Accessed: March 7, 2023.

[81] W3C. 2022. Privacy Community Group. <https://www.w3.org/community/privacycg/>. Accessed: March 7, 2023.

[82] Miranda Wei, Madison Stamos, Sophie Veys, Nathan Reitinger, Justin Goodman, Margot Herman, Dorota Filipczuk, Ben Weinshel, Michelle L. Mazurek, and Blaise Ur. 2020. What Twitter Knows: Characterizing Ad Targeting Practices, User Perceptions, and Ad Explanations Through Users' Own Twitter Data. In *USENIX Security*. USENIX Association, CA, USA, 145–162. <https://www.usenix.org/conference/usenixsecurity20/presentation/wei>

[83] Wikipedia. 2022. Do Not Track. https://en.wikipedia.org/wiki/Do_Not_Track. Accessed: March 7, 2023.

[84] Shitong Zhu, Xunchao Hu, Zhiyun Qian, Zubair Shafiq, and Heng Yin. 2018. Measuring and Disrupting Anti-Adblockers Using Differential Execution Analysis. In *NDSS*. Internet Society, CA, USA, 1–15. <https://doi.org/10.14722/ndss.2018.23331>

[85] Sebastian Zimmeck. 2021. Opting Out May Not Prevent Websites From Collecting Your Data. <https://sebastianzimmeck.medium.com/opting-out-may-not-prevent-websites-from-collecting-your-data-cfc3ff5b5ff7>. Accessed: March 7, 2023.

[86] Sebastian Zimmeck. 2021. Standardizing Global Privacy Control (GPC) #10. <https://github.com/privacycg/proposals/issues/10>. Accessed: March 7, 2023.

[87] Sebastian Zimmeck and Kuba Alicki. 2020. Standardizing and Implementing Do Not Sell. <https://sebastianzimmeck.de/zimmeckAndAlicki2020DoNotSell.pdf>. In *WPES*. ACM, Virtual Event, USA, 1–6.

[88] Sebastian Zimmeck and Steven M. Bellovin. 2014. Privee: An Architecture for Automatically Analyzing Web Privacy Policies. <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/zimmeck>. In *USENIX Security*. USENIX Association, San Diego, CA, USA, 1–16.

A APPENDIX

A.1 Browser Setup Survey Questionnaire

- Would you create an account and log in to sync your data? [Multiple choice; answer required]
 - I would create an account and log in to sync my data
 - I would *not* create an account and log in to sync my data
- Would you enable Global Privacy Control? [Multiple choice; answer required]
 - I would enable Global Privacy Control
 - I would *not* enable Global Privacy Control
- It is important that you pay attention to this study. Please select “Somewhat agree.” [Multiple choice; answer required]
 - Strongly agree
 - Somewhat agree
 - Somewhat disagree
 - Strongly disagree
- Would you like to have a tab bar? [Multiple choice; answer required]
 - I would like to have a tab bar
 - I would *not* like to have a tab bar
- Which theme would you select? [Multiple choice; answer required]
 - System Default
 - Light
 - Dark
- Do you have comments, suggestions, or questions? If so, please let us know. [Long answer text; answer not required]

[Screenshots of the GIFs used in the Browser Setup Survey are shown in Figure 2.]

A.2 GPC Survey Questionnaire

- What is most concerning to you about online privacy, if anything? [Long answer text; answer required]
- In general, are you comfortable with a website showing you ads based on your activity on that website? Assume that the website can collect your data for its own purposes but will not be allowed to share it with advertisers or other companies. [Multiple choice; answer required]
 - Yes
 - No
- Imagine that web browsers, for example, the mobile browser shown below, come with a new privacy feature called Global Privacy Control (GPC). What do you expect to happen if you enable GPC? (assuming you are a California resident) [Multiple choice; answer required]
 - Websites respecting GPC would be prohibited from collecting data from you
 - Websites respecting GPC would be prohibited from giving advertisers your data
 - Websites respecting GPC would be prohibited from showing you advertising
 - GPC signals must be respected for you and anyone else in the world
 - None of the above
 - Other [Short answer text; answer required if selected]

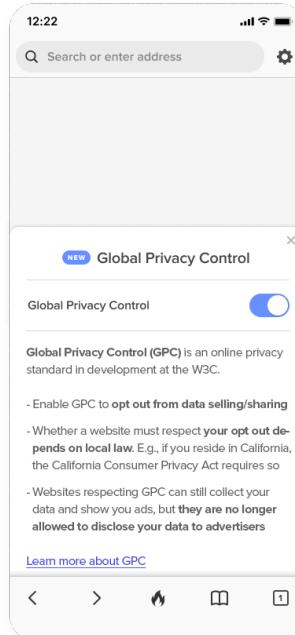


Figure A.1: The screenshot shown to participants to explain GPC and the CCPA's opt out right.

- Please explain in your own words what GPC does. (assuming you are a California resident).
[Long answer text; answer required]
- Would you choose to enable GPC in your browser? Why or why not? (assuming you are a California resident)
[Long answer text; answer required]
- It is important that you pay attention to this study. Please select “Strongly disagree.” [Multiple choice; answer required]
 - Strongly agree
 - Somewhat agree
 - Somewhat disagree
 - Strongly disagree
- To how many websites would you send GPC signals if you could pick them individually? (assuming you are a California resident)
[Multiple choice; answer required]
 - To all websites I visit
 - To most websites I visit but with some exceptions
 - To only a few selected websites I visit
 - To no website I visit
- How satisfied would you be with the Global Privacy Control process in the mobile browser of this survey? (assuming you are a California resident) [Multiple choice; answer required]
 - Very satisfied
 - Somewhat satisfied
 - Somewhat dissatisfied
 - Very dissatisfied
- Do you have comments, suggestions, or questions? If so, please let us know. [Long answer text; answer not required]

A.3 Websites Compliant with GPC

As of August 2022 a total of 27/64 sites of the Combined Set that implement the US Privacy String were compliant and respected GPC.

- al.com
- arstechnica.com
- bloomberg.com
- cnn.com
- cpanel.net
- fandom.com
- fortune.com
- freep.com
- howstuffworks.com
- indiegogo.com
- latimes.com
- mediafire.com
- mlive.com
- nbcnews.com
- newyorker.com
- nj.com
- nydailynews.com
- sciencedaily.com
- slate.com
- theatlantic.com
- theguardian.com/us
- time.com
- usatoday.com
- washingtonpost.com
- weather.com
- webmd.com
- wired.com