AI-Control: Opt-Out Mechanisms From the View of a Governance Cycle

Carl Gahnberg[1]

IAB Workshop on AI-CONTROL, August 2024

Introduction

This paper seeks to inform discussions on an AI-Control mechanism by outlining considerations from a view of governance. Understood as a mechanism for content creators to opt out of having their content used as training data for the creation of large language models (LLMs), an AI-Control mechanism would offer a standard for signaling a content creator's preferences to a web crawler. Starting with a theoretical account of governance and the role of standards in solving strategic problems of coordination and collaboration, the paper introduces the view of a "governance cycle" as a means to identify important considerations in the creation of such an opt-out mechanism. Using this as a framework of analysis, the remaining part of the paper illustrates insights from two previous opt-out mechanisms aimed at similar problem structures, the Do Not Track (DNT) and Global Privacy Control (GPC) initiatives, and their lessons for AI-Control. The paper emphasizes that standards in this context not only facilitate technical interoperability but that viewing their role as part of a broader governance cycle helps identify critical factors for their success.

Governance through standards: Rules for Coordination or Cooperation

The concept of governance, as understood in this paper, is based on Biersteker's definition, as an *"intersubjectively recognized, purposive order […] which defines, constrains and shapes actors' expectations in an issue domain"* [1]. Important in this definition is an understanding of governance as a system of authoritative rules that ultimately come to influence how different actors behave in relation to a given issue. Such rules may be codified and formal, such as laws or regulations, but also encompass informal rules like norms of behavior. Notably, the definition does not necessitate such rules to have been developed by governments but instead emphasizes that rules may be authoritative through a normative belief that they "ought to be obeyed"[2]. Critically, it is a definition based on observations that governance rules in a given domain are sometimes authored by, or with the help of, non-state actors[2].

In this context, standards are seen as governance rules that aim to create order in a given domain. Typically, they are conceptualized as representing solutions to strategic problems of coordination or

[1] This work is performed independently of the author's employer and any author affiliation listed in relation to this document is for identification purposes only. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of their employer.

[2] The notion of a "non-state" or "private" actor is best understood as a residual concept that encompasses any actor that is not a state or a direct agent of a state. Examples include businesses, NGOs, civil society organizations and academia[3].

collaboration and theorized as addressing issues of externalities, understood as a situation in which “one actor’s conduct affects the well-being of another”[4].

Standards aimed at addressing **coordination problems** involve scenarios where actors have a strong incentive to agree on a single rule and have little incentive to defect from the rule once it has been established. This type of strategic situation is commonly associated with developing technical standards for products and processes. The goal is often to reduce transaction costs by devising a single rule for all to use, which results in network effects that quickly raise the cost of unilateral noncompliance. The Internet Protocol (IP) is a classic example of such a situation since the value of being connected to the Internet increases as more users join the network (i.e., adopt the IP standard). This also means that once a standard has been adopted and deployed, all participants have strong incentives to adhere to the standard since doing so aligns with individual benefits.

In contrast, standards seeking to address **collaboration problems** represent scenarios in which an outcome is dependent on actors working together, but at least one party has an incentive to defect rather than cooperate. In such scenarios, the role of a standard is aimed at stipulating the required behaviors and mechanisms to ensure cooperation, thereby helping align objectives and incentives. While adherence to this type of standard may be enforced by laws and regulations [5], or the threat thereof[6], there are also examples of market-based solutions, such as “Fairtrade” labels. By meeting the standard’s criteria and signaling adherence to consumers (via a label), businesses are able to extract a premium price for a product with socially desirable properties. This helps resolve the collective action problem but is dependent on mechanisms for monitoring compliance, which allows consumers to discriminate between those who adhere to the desirable behavior and those who do not[7].

It's important to note that the descriptions of coordination or collaboration problems are theoretical ideals. While they are useful lenses for analyzing a given issue, reality is often more nuanced. This includes defining the “socially desirable outcome” and the fact that some cases may incorporate features from both problems[4]. Furthermore, standards are not independent means of governance but exist within contexts where they interact or link to other rules.

With this in mind, Eberlein[8] describes governance in a given domain as a process that can be separated into six separate yet complementary steps, known as a “governance cycle”. Such perspective can also be applied to the role of a standard, in which the standardization process can be understood as a distinct step within a larger process, and as linked to other rules and mechanisms in a broader governance framework:

1. **Framing the regulatory agenda and setting objectives:** The first step of the governance cycle involves framing the problem to be addressed and outlining the objectives of the new rules. It’s also at this stage that, after articulating a particular problem, the rights and responsibilities of different actors are identified and/or debated.
2. **Formulating rules or norms:** The core of any governance effort are the rules (formal or informal) that define the expected behaviors. What those are is typically linked to objectives identified in the regulatory agenda since prescribed obligations and behaviors typically seek to align with identified rights and responsibilities. In the context of an opt-out mechanism, this step would encompass the standardization of the technical mechanism. Defining the technical specification can thus be understood as formulating a procedural rule for how to communicate substantive rules that invoke behavioral demands from the counterpart.
3. **Implementing rules within targets:** The development of new rules does not guarantee adoption. This is well understood in processes of technical standardization, where the development of a specification or progressing its status along a formal track does not guarantee its use.

4. **Gathering information and monitoring behavior:** Confirming whether actors comply with agreed rules may be required, often dependent on the problem a governance rule seeks to address. For example, standards deployed in coordination problems typically do not require monitoring since adhering to the standard aligns with individual benefits. In contrast, cooperation problems are likely to demand greater monitoring since one or more actors may have an incentive to deviate from the standard.
5. **Responding to non-compliance via sanctions and other forms of enforcement:** What actions are required in response to non-compliance may vary based on the type of problem. In coordination issues, consequences are often direct and clear, such as exclusion from a market due to non-adherence to technical standards. In cooperation problems, sanctions may require the imposition of reputational or monetary costs or even legal sanctions.
6. **Evaluating policy and providing feedback, including reviewing rules:** Finally, governance processes are typically assessed on the effectiveness of rules and the overall governance framework to inform future improvements. This involves reviewing whether initial objectives were met, identifying unintended consequences, and updating rules to remain relevant and effective.

The governance cycle is a conceptual ideal but helps highlight that the effectiveness of a given governance rule may depend on interactions with a broader set of rules, and where different actors may contribute or participate in different stages of the process. For example, activists can promote ethical principles that shape regulatory objectives, businesses may adopt standards developed by competitors, and public actors can help support enforcement. Importantly, it also helps illustrate that a governance process can have dependencies across steps, and the importance of considering the implications for subsequent phases of the cycle.

AI-Control: A Collaboration Problem in Need of Coordination

The creation of an AI-Control mechanism, understood as a mechanism for content creators to opt out of having their content used as training data for the creation of large language models (LLMs), seeks to address the overarching problem of signaling a content creator's preferences to a web crawler. However, using the framework outlined above, the creation of such a mechanism can also be seen as a mix of two strategic problems: one of collaboration and one of coordination.

First, an AI-Control mechanism would be fundamentally aimed at addressing the problem of collaboration. From this perspective, an AI-Control mechanism seeks to address a problem in which the socially desirable outcome is one where content creators can assert their applicable rights. However, due to the value of using the content in training AI models, at least one party, the AI-model creators, has a strong incentive to defect from the collaboration – understood as not respecting the rights of the content creator. It's important to note that such defection is not guaranteed but effectively a cost-benefit calculation: if the benefits of violating the content creator's rights outweigh the cost, then this is likely to happen. For example, being caught violating the applicable rights may have reputational costs to the AI creator, which might outweigh the benefits of using the content. However, if the value of the content is sufficiently high, or the reputational cost is low or non-existent (no one notices the violation), then defection is likely.

While the content creator's assertion of their rights is a substantive rule addressing a collaboration problem by invoking behavioral demands from the AI creator there is also the need for a technical mechanism to communicate the demand. This can be thought of as a procedural rule aimed at addressing a problem of coordination. Basically, given that coordination is demanded[3] by all parties they would need to develop a technical standard that is able to instruct how an automated crawler should treat a given site.

As discussed above, the separation between collaboration and coordination problems should be understood as theoretical ideals, and reality may be more nuanced and complex. Consider the technical mechanism that, although framed as addressing a problem of coordination, may also impact the collaboration problem. For example, if the signal is ambiguous, or generally unreliable, this could be used as an excuse by AI creators seeking to defect from the substantive rules. In other words, while theoretically separate, there may be "leaks" across the two problem structures that make reality much more complex.

DNT and GPC – lessons from a view of the governance cycle

To illustrate considerations for an AI-Control mechanism from the perspective of a governance cycle, this paper presents a brief case study of two previous opt-out mechanisms: the Do Not Track (DNT) effort and the Global Privacy Control (GPC). These were selected since both allow users to signal their preferences regarding the collection of their data by third parties but have had different outcomes; importantly, the GPC has seen legal enforcement, unlike DNT.

The experiences with DNT and GPC provide useful insights for discussions on developing a new AI-Control mechanism, notably because they address a similar problem structure of collaboration. However, several key differences must be considered when applying these lessons to an AI-Control mechanism. First, the roles of the parties involved are reversed. In DNT and GPC, the client (user) signals their preferences to a server. In contrast, in the scenario of an AI-Control mechanism, it would be the server that signals preferences or restrictions to an (automated) client, such as a web crawler. Secondly, the DNT and GPC protocols utilize binary signals (on/off) to indicate an opt-out preference. However, an AI-Control mechanism may require more nuanced signaling options. For instance, content creators might want to specify specific use policies depending on how the trained model will be used (e.g., for commercial or non-commercial purposes).

While the list of differences can be made longer (including some further illustrated below) there are lessons from the two cases that can be applied in developing an AI-Control mechanism. To this end, the following sections uses the six steps of a governance cycle to illustrate insights from DNT and GPC that could help inform discussions on AI-Control.

1. Framing the Regulatory Agenda and Setting Objectives

The case of DNT offers a useful illustration of how ambiguity in a regulatory agenda can create fundamental challenges in subsequent stages of the process. While motivated by privacy concerns related to online tracking, the DNT initiative was not anchored in privacy legislation. Instead, it was launched with the idea that new legislation would not be necessary, as existing consumer protection regulations could enforce it. Specifically, the Federal Trade Commission's (FTC) powers under Section 5 to act against deceptive trade practices were considered sufficient, where deviations from a self-regulatory code of conduct on 'Do-Not-Track' would be addressed.

An important implication of this framing was that the very definition of 'tracking,' and consequently the governance objectives, were contested from the outset. Civil society organizations, prioritizing

[3] Note that the demand for coordination may stem from coercion. For example, an AI creator may face legal sanctions if they disregard or fail to acknowledge an opt-out signal and thus incentivized to find an efficient coordination mechanism that minimize such risks.

privacy, advocated for a DNT interpretation focused on limiting data collection. Conversely, the advertising industry viewed the DNT signal as a mechanism for users to opt out of targeted advertising while still permitting data collection. Although most stakeholders perceived the issue primarily as one of data collection, the FTC's inability to clearly define consumer rights in this context led to ambiguous objectives. It was unclear whether the signal was intended to invoke a right to opt out of data collection entirely or merely from targeted advertising, resulting in unclear obligations for the advertising industry.

In stark contrast, the Global Privacy Control (GPC) was developed in response to California's Consumer Privacy Act (CCPA), which explicitly grants users the right to opt out of the sale of their personal information[9]. This meant that the GPC mechanism could be based on clearly defined rights (users can object to the sale of their data to third parties) and obligations (companies must not sell a user's data if requested not to). An important implication of this is that the CCPA recognized that the collection of a user's data by a third party depended on the first party's action of 'selling' this data. In other words, by framing the issue around the sale of users' data, the CCPA not only defined clear rights and obligations but also shifted the responsibility of respecting a user’s rights to the first party.

Lessons for AI control: The cases of DNT and GPC illustrate the role of clearly defined rights and responsibilities in defining the substantive rules of an opt-out mechanism. In the case of GPC this was clearly aided by underlying legislation offering these definitions, while DNT was initiated on a highly contested understanding of the problem to be resolved.

The case for an AI-Control mechanism has a relatively strong starting point in this light since it links to a well-established regime for intellectual property rights. As described by Nottingham[10], this includes existing provisions that indicate that the use of an opt-out signal could be used by a content creator to assert their rights. However, current debates also point to these rights being contested in the context of training AI models, as illustrated by prominent legal disputes in the United States. This implies that the creation of an AI-Control mechanism will have a dependency on how existing rights are to be interpreted and/or updated in the context of training AI models.

2. *Formulating Rules or Norms*

The distinction between procedural and substantive rules is evident in the history of DNT, for which the standardization effort was divided into two tracks: 1) the 'Tracking Preference Expression' (TPE), which defined the technical mechanisms for expressing a Do-Not-Track preference; and 2) the 'Tracking Compliance and Scope' (TCS), which defined the meaning of a Do-Not-Track preference and the practices required by websites to comply.

Work on the TPE was relatively straightforward, as the approach of using HTTP headers had already been deployed by large browser vendors. However, the TCS, which addressed the meaning of the DNT signal, proved more contentious. For example, the discussion became focused on the very meaning of “tracking” and identifying boundaries for what data could be collected for other permitted uses, such as security. As mentioned earlier, this was partly due to the ambiguity in the regulatory agenda regarding which rights were to be respected. Consequently, the DNT effort failed to reach a consensus on its substantive rules.

In contrast, while the GPC uses a technical mechanism based on DNT's TPE, its substantive rules are based on and defer to legislation to specify the required behavior (i.e., 'do not sell'). Notably, just like DNT, the substantive part of GPC is coded in a binary signal that invokes behavioral demands from the counterpart, which are specified elsewhere[4]. Yet, while DNT sought to specify those behavioral demands in TCS, GPC deferred those behavioral demands to laws and regulations by specifying that the GCP signal is *"[…] a person's assertion of their applicable rights to prevent the sale of their data, the sharing of their data with third-parties, and the use of their data for cross-site targeted advertising"[12].*

Lesson for AI control: The case of GPC clearly illustrates how the substantive rule in an opt-out mechanism can benefit from deferring the behavioral demands to legislation. This is particularly important in a scenario where an AI-Control mechanism goes beyond a binary signal, and towards more granular demands of behaviors. It also illustrates how the development of a new mechanism must seek input from a broader IPR regime, to which a signal could defer the behavioral demands - whether from legislation or new licensing schemes.

Furthermore, the two cases also illustrate the importance of considering the target of a signal, and how the substantive rule would be articulated. For example, whether the substantive rule is directly addressed to the web crawler and restricts the collection of content for certain ends, or if the signal targets the AI creator, restricting the use of the content for certain ends.

3. Implementing Rules Within Targets

The beginnings of the DNT effort were highly successful in attracting all of the relevant parties to work on the standardization effort. This was in part due to regulatory pressure and expectations of legislative actions at the time [13] but also because of early implementations of the DNT signal by many of the major browsers[14]. In fact, DNT adoption has, from a browser perspective, remained relatively high, with most of the larger vendors still offering users the ability to send a DNT signal. Instead, the critical challenge for DNT has been the adoption on the server side, with few websites honoring the signal. This lack of adoption, and in the absence of legal requirements, is fundamentally linked to a lack of consensus on the very meaning of a DNT signal (i.e., its substantive rules) and, by extension, the behavioral demands.

However, one of the key controversies in the early days of the effort was also about the user's intent and specifically about having DNT as the default setting. For some privacy advocates setting DNT to default would reflect users' preference for privacy by default, while opponents argued that the self-regulatory approach necessitated a consumer choice to activate DNT. The debate culminated in 2012 when Microsoft announced that its IE 10 would come with DNT set as the default, provoking a strong response from the Digital Advertising Alliance (DAA), stating that its members could ignore DNT signals from IE 10[15].

The GPC has seen relatively low implementation amongst major browsers but is available as an add-on to most. On the server side, however, a recent study [11] found that many sites did not respect the GPC signal, with only 12% indicating compliance. With regard to user adoption, the GPC also notes that the validity of a signal set by default may vary between jurisdictions. For example, a signal set by default may be perfectly valid in those where users are required to opt-in to the sale of their data, whereas others may require the user to actively turn on the GPC signal for it to be considered a valid expression of intent to opt-out[11]. This again illustrates the GPC's approach of deferring discussions on the substantive rules to the relevant jurisdiction.

Lessons for AI control: In both DNT and GPC, the approach of implementing the signal through the HTTP header proved relatively easy. However, as seen by the server-side adoption rate, this is only half of the story. The case of an AI-Control mechanism is likely to look quite different,

[4] Aside from offering simplicity, the use of a binary signal also offers particular benefits in the case of privacy-related opt-out mechanisms since it reduces the risk of "fingerprinting" [11]

depending on how it is implemented. For example, if an AI-Control mechanism builds on robots.txt, the procedural rule would, in effect, already be in place, allowing standardization efforts to focus on the substantive rule to be communicated. In contrast, a new mechanism may face greater challenges for wider implementation since, unlike the cases of GPC and DNT's browser implementations, there are no intermediary authorities to spur the adoption rate.

4. Gathering Information and Monitoring Behavior

The DNT effort has been criticized from its very beginnings due to the difficulties of monitoring compliance by third parties embedded in a website. This can be illustrated by the Electronic Frontier Foundation's (EFF) effort to develop its own policy as an alternative to the TCS. In contrast to the TCS, and due to the inability of monitoring compliance by third parties, EFF's policy relied on the first party to get a contractual commitment from its embedded third parties to respect DNT - effectively outsourcing the monitoring task to the first party[16].

This shift in focus towards the first party is inherent to the GPC signal since it targets the sale of a user's data. Since the signal targets the first party, and since it's sent before the loading of the page, it allows the first party to block scripts of a third party when responding to the user's request. This means that compliance with the signal can also be more easily monitored since the loading of third party scripts can be audited. For example, in its enforcement of the GPC signal the California Attorney Genereal (AG) described how the investigation used commercially available browser extensions to monitor data flows to third-parties, which allowed the investigators to compare traffic patterns when activating the GPC signal[17].

Lessons for AI control: From the experiences of DNT and GPC it is clear that the ability to monitor compliance is crucial for the success of an opt-out mechanism. In the case of the AG's investigation of GPC, this ability was in part a consequence of the target being the first party, but also due to the signal being binary, which simplified the interpretation of a behavioral change.

Given that an AI-Control mechanism seeks to address a problem of collaboration, the ability to monitor compliance will be critical to its success. And while an implementation using robot.txt could support such effort through investigation of server logs, it might be important to consider if this is an efficient solution, notably from a view of scalability and third-party verification (e.g., by law enforcement agencies doing broad enforcement sweeps, as in the case of GPC) .

5. Responding to Non-Compliance via Sanctions and Other Forms of Enforcement

While the DNT initiative has seen efforts to enforce a signal through technical means, such as via EFFs Privacy Badger, there has not been a legal enforcement of DNT. While this lack of enforcement in the US can be explained by a lack of consensus on TCS, which prevented enforcement by the FTC, the European context was also hampered by institutional constraints. For example, the Article 29 Working Party, composed of members from all of the EU member states' Data Protection Authorities (DPAs), had been actively monitoring the DNT process and provided feedback on requirements for using DNT in a European context[18]. Yet, while the Article 29 WP could have provided guidance on the enforcement of DNT they were unlikely to do so for a specification that was only at the earlier stages of the W3C's standards track.

The GPC, on the other hand, has been enforced. A significant milestone was reached in August 2022 when the California AG settled with the French online retailer Sephora for failing to honor GPC opt-out requests (Bonta, 2022a). It illustrated that enforcement ultimately relies on the commitment of public institutions. This commitment was evident in public statements by the

Attorney General affirming that GPC is a valid opt-out mechanism (REF). Moreover, the appointment of Ashkan Soltani, former Chief Technology Officer at the FTC and a contributor to the GPC's development, as the Executive Director of the California Privacy Protection Agency (CPPA) further demonstrated a clear intent to enforce the signal (LINC, 2021).

Lessons for AI-Control: Both DNT and GPC illustrate the importance of close collaborations with public authorities that could enforce an opt-out signal. While institutional constraints in the EU impeded the collaboration on DNT, the GPC clearly illustrated how the capacity and willingness to enforce may be crucial for success.

To this end, there is a need to actively include public authorities in the work on an AI-Control mechanism. Both to ensure that its substantive rule is able to invoke legal rights and obligations, but also to ensure that enforcement authorities have the capacity to perform the investigations. This includes working with multiple jurisdictions and identifying any institutional constraints that may impede a public authority's ability to enforce (as in the case of WP29 and progression along a standards track).

6. Evaluating Policy and Providing Feedback, including Review of Rules

While DNT has ultimately failed in its efforts to govern online tracking, there have been a number of attempts to evaluate and improve it. The EFF's efforts to promote its own DNT policy, as described above, is one such example in which stakeholders sought to address failures identified in the process. Other examples include efforts to adapt the DNT's TPE following the EU's adoption of the GDPR and negotiations on a new e-Privacy Regulation (ePR)[19].

This feedback loop is also visible in the GPC, which has relied extensively on the DNT experience. The fact that the GPC's technical mechanism is fundamentally based on TPE is an obvious example, but insights gained from DNT have also been more substantive. For example, the GPC was in part a response to public comments by the California AG on the use of DNT as an opt-out mechanism under the CCPA regulations, in which the AG notes that the signal is understood as a request to third-parties and that a new signal would be needed to clearly express a user's intent to opt-out of the sale of personal information[20].

Lesson for AI-Control: While the description of a governance cycle tend to present a governance effort as having a clear start and finish, neatly fitted along the six steps, reality may look different. For example, in the case of GPC the effort did not start from scratch but could rely on DNT's work of specifying a technical mechanism of TPE. Furthermore, GPC had the benefit of learning from the DNT effort's failures, and to adapt accordingly.

In a similar vein, an AI-Control mechanism can learn from previous efforts. This includes continued research on the use of robots.txt and associated behaviors, but also other efforts to have clients retrieve server-side policies. Importantly those that share a similar problem strucutre and that might have failed (e.g. the Platform for Privacy Preferences Project, P3P).

Conclusion

To illustrate considerations for designing an AI-Control mechanism, this paper has described how standards function as governance tools and as means to resolve strategic problems of collaboration or coordination. Through the lens of a governance cycle, it has presented lessons from the DNT and GPC initiatives, which represent two opt-out mechanisms designed for similar problem structures. As illustrated by the two cases, the design of an AI-Control mechanism needs to consider the importance of clear rights and responsibilities, and the value of anchoring substantive rules in a legislative framework. Furthermore, the two cases also illustrate how a mechanism's design and the signals target may influence important abilities of monitoring compliance. Finally, given that an AI-Control mechanism seeks to address a problem of collaboration, and as illustrated by DNT and GPC, there is a need actively engage with public authorities to support interpretation and enforcement of an opt-out mechanisms. Future work on an AI-Control mechanism can benefit from these insights to help structure discussions and foster a more efficient use of standards within broader governance frameworks.

Bibliography

[1] T. J. Biersteker, 'Global governance', in *The Routledge Handbook of Security Studies*, M. Dunn Cavelty and V. Mauer, Eds., London: Routledge, 2010, pp. 455–467.

[2] I. Hurd, 'Legitimacy and authority in international politics', *Int. Organ.*, vol. 53, no. 02, pp. 379–408, 1999.

[3] B. Arts, 'Non-state actors in global governance: three faces of power', Preprints aus der Max-Planck-Projektgruppe Recht der Gemeinschaftsgüter, 2003.

[4] K. W. Abbott and D. Snidal, 'International'standards' and international governance', *J. Eur. Public Policy*, vol. 8, no. 3, pp. 345–370, 2001.

[5] T. Büthe and W. Mattli, *The new global rulers: The privatization of regulation in the world economy*. Princeton University Press, 2011.

[6] P. Verbruggen, 'Gorillas in the closet? Public and private actors in the enforcement of transnational private regulation', *Regul. Gov.*, vol. 7, no. 4, pp. 512–532, 2013.

[7] C. Roger and P. Dauvergne, 'The rise of transnational governance as a field of study', *Int. Stud. Rev.*, vol. 18, no. 3, pp. 415–437, 2016.

[8] B. Eberlein, K. W. Abbott, J. Black, E. Meidinger, and S. Wood, 'Transnational business governance interactions: Conceptualization and framework for analysis', *Regul. Gov.*, vol. 8, no. 1, pp. 1–21, 2014.

[9] GPC, 'Global Privacy Control', Global Privacy Control. [Online]. Available: https://globalprivacycontrol.org/

[10] M. Nottingham, Considerations for AI Opt-Out. Accessed: Aug. 02, 2024. [Online]. Available: https://www.mnot.net/blog/2024/04/21/ai-control

[11] S. Zimmeck, O. Wang, K. Alicki, J. Wang, and S. Eng, 'Usability and enforceability of global privacy control', *Proc. Priv. Enhancing Technol.*, vol. 2023, no. 2, 2023.

[12] S. Zimmeck, P. Snyder, J. Brookman, and A. Zucker-Scarff, *Global Privacy Control (GPC): Proposal 20 April 2023*, Apr. 20, 2023. Accessed: Apr. 23, 2023. [Online]. Available: https://privacycg.github.io/gpc-spec/

[13] A. McDonald, 'Stakeholders and High Stakes: Divergent Standards for Do Not Track', *Camb. Handb. Consum. Priv. Camb.*, 2018.

[14] R. Singel, 'Burning Question: Should I Use My Browser's Do-Not-Track Setting?', *WIRED*, Jul. 26, 2011. [Online]. Available: https://www.wired.com/2011/07/pr-burning-donottrack/

[15] DAA, 'Digital Advertising Alliance Gives Guidance to Marketers for Microsoft IE10 "Do Not Track" Default Setting'. [Online]. Available: https://digitaladvertisingalliance.org/press-release/digital-advertising-alliance-gives-guidance-marketers-microsoft-ie10-%E2%80%98do-not-track%E2%80%99

[16] Electronic Frontier Foundation (EFF), 'Understanding EFF's Do Not Track Policy: A Universal Opt-Out From Tracking', Understanding EFF's Do Not Track Policy: A Universal Opt-Out From Tracking. [Online]. Available: https://www.eff.org/pages/understanding-effs-do-not-track-policy-universal-opt-out-tracking

[17] A. G. Bonta, 'PEOPLE OF THE STATE OF CALIFORNIA v. SEPHORA USA, INC. - COMPLAINT FOR INJUNCTION, CIVIL PENALTIES, AND OTHER EQUITABLE RELIEF'. Aug. 23, 2022. Accessed: Apr. 23, 2023. [Online]. Available: https://oag.ca.gov/system/files/attachments/press-docs/Complaint%20%288-23-22%20FINAL%29.pdf

[18] WP29, 'Article 29 Data Protection Working Party comments in response to W3C's public consultation on the W3C Last Call Working Draft, 24 April 2014, Tracking Preference Expression (DNT)', Jun. 06, 2014. Accessed: Apr. 23, 2023. [Online]. Available: https://lists.w3.org/Archives/Public/public-tracking-comments/2014Jun/0000.html

[19] M. Schunter, *Tracking Protection Working Group Charter*, Dec. 31, 2016. [Online]. Available: https://www.w3.org/2016/11/tracking-protection-wg.html

[20] A. G. Becerra, 'FSOR APPENDIX E: SUMMARY AND RESPONSE TO COMMENTS SUBMITTED DURING SECOND 15-DAY COMMENT PERIOD'. State of California Department of Justice, Jun. 01, 2020. Accessed: Apr. 23, 2023. [Online]. Available: https://oag.ca.gov/privacy/ccpa/regs