University Library



A gateway to Melbourne's research publications

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:
Dreyfus, S;Tashevtseva, N;Parampalli, U

Title:
The Harm Frontier in Practice: A Case Study of Privacy Tracking

Date:
2025

Citation:
Dreyfus, S., Tashevtseva, N. & Parampalli, U. (2025). The Harm Frontier in Practice: A Case Study of Privacy Tracking. ACIS 2025 Proceedings, pp.1-20. AIS.

Persistent Link:
https://hdl.handle.net/11343/369020

License:
CC BY-NC

Association for Information Systems
AIS Electronic Library (AISeL)

ACIS 2025 Proceedings | Australasian (ACIS)

12-1-2025

The Harm Frontier in Practice: A Case Study of Privacy Tracking

Nataliia Tashevtseva
The University of Melbourne, natasha.tashevtseva@unimelb.edu.au

Suelette Dreyfus
The University of Melbourne, suelette@unimelb.edu.au

Udaya Parampalli
The University of Melbourne, udaya@unimelb.edu.au

Follow this and additional works at: https://aisel.aisnet.org/acis2025

Recommended Citation

Tashevtseva, Nataliia; Dreyfus, Suelette; and Parampalli, Udaya, "The Harm Frontier in Practice: A Case Study of Privacy Tracking" (2025). *ACIS 2025 Proceedings*. 295.
https://aisel.aisnet.org/acis2025/295

This material is brought to you by the Australasian (ACIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ACIS 2025 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

The Harm Frontier in Practice: A Case Study of Privacy Tracking

Full research paper

Nataliia Tashevtseva
School of Computing and Information Systems
The University of Melbourne
Melbourne, Australia
Email: nstrelnikova@student.unimelb.edu.au

Suelette Dreyfus
School of Computing and Information Systems
The University of Melbourne
Melbourne, Australia
Email: suelette@unimelb.edu.au

Udaya Parampalli
School of Computing and Information Systems
The University of Melbourne
Melbourne, Australia
Email: udaya@unimelb.edu.au

Abstract

This paper introduces the concept of the Harm Frontier for AI impact on privacy to describe the expanding gap between technical capabilities for privacy invasion and the regulatory mechanisms meant to control them. Through an audit of Australia's most popular websites and an in-depth case study of a major news outlet, the paper offers a practical example of the Harm Frontier case. We also discuss how the widespread inclusion of AI in digital systems is widening that gap and how the indirect use of AI technologies, combined with underregulated and non-transparent data collection, is accelerating the growth of the Harm Frontier.

Keywords Artificial Intelligence (AI), privacy, regulatory gap, website tracking, harm frontier

1 Introduction

As digital technologies evolve at an unprecedented pace and interconnect every aspect of society, the protection of individual privacy has become a critical concern (Oluwatosin Reis et al. 2024; Romansky and Noninska 2020). Among these technologies, Artificial Intelligence (AI) stands out not only for its capacity to automate tasks and optimise systems but also for its role in reshaping the privacy landscape (Ashok et al. 2022; Dwivedi et al. 2021). AI technologies are increasingly embedded in consumer-facing applications across domains such as e-commerce, social media, digital health, entertainment platforms, and personalised finance. These systems benefit both consumers and organisations, yet these conveniences come at a privacy and autonomy cost. The widespread use of AI for the collection, analysis and inference of user data raises pressing ethical and privacy concerns.

The black box nature of AI systems, the lack of transparency of their use and the absence of meaningful user control or consent mechanisms undermine core privacy principles. Moreover, AI's capacity to make inferences and decisions without human oversight enables continuous behavioural surveillance often without the user's awareness or understanding, expanding the possibility of wrong perceptions, discrimination and ethical issues. AI inclusion in the user data analysis creates a breach in the social norms (Marks 2021). Users are often unaware of online data collection practices: which data is collected or how this data is further processed and repurposed (ACCC 2024),(Andreotta et al. 2022), the inference practices and the extent of tracking, especially when AI systems are implicit in digital interactions. In the Digital Platform Services Inquiry report, the ACCC highlights that this lack of awareness seriously undermines consumers' ability to make informed decisions or consent (ACCC 2024). AI inclusion further exaggerates this by providing more opportunities for data inference, large amounts of data processing and the fact that many AI systems operate invisibly, collecting behavioural and contextual data far beyond what users knowingly consent to. AI is increasingly deployed within websites directly or through third parties. Yet there are no current legal requirements for these systems to be identified to users. While the Office of the Australian Information Commissioner (OAIC) encourages transparency, there's no legal standard for private companies to explicitly disclose the use of AI technologies on their websites for data collection or processing purposes. Profiling or inference is often masked under vague terms like "analytics," "optimisation," "experience personalisation", or "service improvement". This regulatory ambiguity leaves users unaware of how AI shapes their experience, while companies operate and profit with minimal accountability, leaving consumers prone to significant privacy risks.

This widening disparity between what is technically possible and what is meaningfully disclosed or regulated forms the basis of what we describe in this paper as the Harm Frontier. The Harm Frontier captures the growing gap between the capabilities of AI systems to collect, infer, and act on personal data, and the regulatory or user-facing mechanisms designed to govern them. While technological change is accelerating, driven by the rapid evolvement of AI-driven tracking, behavioural analytics, and real-time personalisation, regulatory response remains slow and fragmented. This gap creates a grey zone in which organisations can deploy increasingly intrusive data practices with minimal transparency or legal obligations, exposing users to privacy risks they cannot anticipate or meaningfully resist.

In the following sections, we introduce the Harm Frontier as a theoretical concept and the Harm Frontier Matrix as a conceptual tool to map this misalignment and identify systems that fall into this regulatory blind spot. While our empirical study does not aim to validate the Harm Frontier model quantitatively, it offers a practical example of the Harm Frontier cases through a privacy audit of Australia's most popular websites, an in-depth case study of a major news outlet, which both illustrate scenarios where AI-enabling data practices are present, but regulation does not clearly address or mitigate resulting risks, thereby situating them along the Harm Frontier. This paper is done as a part of the thesis research, which aims to create a Framework for measuring privacy in AI-driven digital systems through a user-centric approach and further contributes to the creation and adoption of ethically aligned AI systems with a focus on preserving privacy and a deeper understanding of how to protect users in systems shaped by automated decision-making. Therefore, it aims to contribute to both academic and industry practice.

2 Background

Artificial Intelligence (AI) is a broad term with no generally accepted definition of the concept (Russell and Norvig, 2020). In this work, we refer to AI as systems and techniques capable of performing tasks that traditionally require human intelligence, such as reasoning, pattern recognition, decision-making, making predictions and language understanding. These systems include, but are not limited to, machine learning (ML) algorithms, natural language processing models, recommender systems, and behavioural inference engines. While we particularly focus on systems that analyse user metadata and behaviour to

make inferences or decisions, we suggest that the Harm Frontier (Section 3) may also extend to other domains, such as generative AI (GenAI) or even generalisation to other technologies, but needs further exploration. The integration of Artificial Intelligence (AI) into consumer digital services has accelerated the process of datafication, which refers to the process of transformation of aspects of human life and social realities into quantifiable digital data, often for purposes of analysis, prediction, and economic value creation (Mejias and Couldry 2019; Kollnig and Shadbolt 2022). This process causes a profound epistemological shift in how knowledge is produced and understood (Chia 2023), turning everyday behaviours into measurable inputs and recasting individuals as sources of monetizable data flows. AI-driven systems converted personal data into a key economic asset that fuels predictive models, optimisation engines, and targeted advertising economies.

2.1 Privacy Risks Amplified by AI

In this study, we look at privacy as the right "to be free from interference and intrusion" and "be able to control who can see or use information about you" as per the OAIC definition, and therefore we understand privacy risk as the potential for systems to make inferences or decisions about individuals using their data and metadata, that go beyond the individual's awareness or consent. These risks are amplified when opaque AI systems are involved, as users are often unaware of the extent or purpose of inference, and those systems have greater capabilities for high-volume inference, personalisation, and behavioural targeting at scale. Rather than treating data collection as a standalone concern, we focus on how such data is repurposed for behavioural prediction, profiling, and surveillance.

Unlike traditional data collection practices, AI operates subtly, integrated into everyday use technologies, silently profiling users, gathering metadata, predicting behaviours, making inferences, and ultimately impacting people's lives. AI systems often operate without users' knowledge or explicit consent (Andreotta et al. 2022). Studies have shown that consumers struggle to understand what data AI collects or how it is used, raising profound concerns around autonomy, informed consent (ACCC 2024) and potential user manipulation (Andreotta et al. 2022; Carroll et al. 2023). As AI continues to expand its reach, this pervasive datafication reshapes both visibility and power in digital consumer environments. AI can be covertly integrated into existing technologies through automatic updates rather than conscious user adoption, affecting billions of users and making it hard for users to detect and evaluate, with potential social and ethical implications (Engström and Strimling 2020). Automated data-gathering systems often operate without meaningful user consent or control; therefore, users have little or no input or awareness of how their data is collected or profiled (Elliott and Soifer 2022).

This pervasive deployment threatens equality, respect, and autonomy, as users often lack real choice regarding AI exposure when using popular platforms. AI is often not seen as something harmful by users due to the lack of awareness, and companies are actively backing up this view by not exposing the actual privacy risks of AI-driven systems. While "traditional" methods of online tracking and data collection already bring significant privacy risks and invasion, AI intensifies this privacy intrusion. Third-party tracking allows companies to collect users' behavioural data, track their activity across digital devices, and potentially share this data with third-party companies (Kollnig and Shadbolt 2022). AI and Machine learning algorithms, in the meantime, can extract behaviours, preferences, and feelings from big data analytics, automatically generating knowledge that forms business decisions but may result in discriminatory perceptions of individuals' lives (Battaglini and Rasmussen 2019). As AI systems operate (collect and process) on vast amount of personal data and without consumer consent it creates situations when AI systems may "know" more about a person than the person knows about themselves, effectively undermining epistemic privilege (Elliott and Soifer 2022).

Furthermore, AI can infer health information from seemingly unrelated digital traces like location data, retail purchases, and social media activity, creating "emergent medical data" without people's knowledge or consent (Marks, 2021), which violates social norms through the erosion of contextual integrity (Helen Nissenbaum's theory of contextual integrity)(Marks 2021). AI models can infer highly sensitive personal information such as race, gender, intelligence, truthfulness, mental state and emotions from seemingly innocuous data sources that humans cannot easily analyse, and those models are becoming widespread (Jurewicz et al. n.d.). Some ML algorithms can now identify depressed users based on Instagram photos or Twitter texts with greater precision than unassisted general practitioners (Reece and Danforth 2017; Reece et al. 2017). Target not only used predictive analytics to infer a teenage girl's pregnancy based on her purchasing habits before even her parents knew, but openly claims that they manipulate their customers by hiding targeted ads between others, "so the baby ads looked random" (Kashmir Hill 2012). TikTok, Instagram, and Facebook routinely infer users' gender, interests and behavioural information based on users' activities (TikTok Privacy Centre n.d.). Facebook claimed to over-use inferred sensitive traits which users have not willingly disclosed through the case-study of

young woman Joanna who complained she had been exposed to content that she believed exploited her mental vulnerabilities (Panoptykon Foundation 2021). There is a growing concern around Femtech applications, period tracking applications in particular, as more studies show that those applications sell and overuse data of their users (Lieber, 2025). In the post-Roe era, users are concerned that Femtech applications with data from period trackers and fertility apps are potentially accessible to law enforcement, advertisers, and third parties - without users' knowledge or meaningful consent (Abigail Dubiniecki 2024). Despite those applications being marked as "health and wellness apps", but the data they collect is not subject to the privacy protections that apply to common medical data (Abigail Dubiniecki 2024).

The recent case study (Lieber, 2025) once again highlights how AI can make high-stakes inferences about individuals without their awareness. Ron Lieber reveals how colleges determine financial aid and pricing for students through complex algorithms developed by consulting firms, such as EAB and Ruffalo Noel Levitz (RNL), which are owned by private equity firms. These algorithms use extensive data about applicants, including demographic, socioeconomic, geographic, academic, curricular, and extracurricular information collected from applications and interactions with the school's digital platforms. First, this process starts with targeted advertisements to these students, tracking their real-time digital behaviour such as email opens, webpage clicks, and time spent on scholarship pages. Then this behavioural data, described as "footprints in the sand," which contains up to 200 variables, is used to assess a student's interest level and willingness to pay. Colleges use data to match students' ability to pay and academic achievements, offering tailored financial aid that is often negotiated before deadlines to maximise enrolment and revenue. This process operates with little transparency, with proprietary algorithms and data practices that families rarely fully understand or can challenge.

2.2 Regulatory Gaps and User Awareness

Current data protection laws contain loopholes that fail to adequately protect users from these automated profiling practices. Current research reveals significant limitations in AI regulation's ability to protect users effectively. Legislators and regulators often lack the technical expertise needed to translate public values into enforceable requirements, while overreliance on industry self-regulation leaves users vulnerable to opaque and unaccountable AI practices (Hadfield and Clark 2023). Traditional privacy models and regulatory tools are poorly equipped to capture the indirect and often opaque ways in which AI shapes user experiences, decision-making, and personal data flows (Lee et al. 2024), (Solow-Niederman 2022), (Martin and Zimmermann 2024).

AI disrupts traditional privacy decision-making frameworks as data ownership becomes blurred, and the traditional risk vs benefits trade-off is not applicable as AI automates data collection so that data disclosure is passive or unnoticed. Users cannot fully assess risks in a traditional way as other factors impact their decision process. Current regulations and tools do not address the layered complexity of AI privacy threats (Martin and Zimmermann 2024). The European Parliament's analysis indicates that the GDPR's principles, like purpose limitation and data minimisation, struggle to align with AI's data-intensive nature, leading to enforcement difficulties (Sartor and Lagioia 2020). An article from the Cambridge Forum on AI points out that existing data protection laws, including the GDPR, face structural challenges when applied to generative AI, often resulting in limited scope and effectiveness. While some technical and policy-based responses have been proposed, the works discussed highlight that those initiatives often remain conceptual, limited in scope, or inaccessible to the average user (Hadfield and Clark 2023; Andreotta et al. 2022). Existing methods, tools, and resources available for AI practitioners are not effective against unique privacy risks introduced or exacerbated by their use of AI in their products (Lee et al. 2024). This combination of technological opacity, limited user awareness, and slow regulatory adaptation sets the stage for what this paper terms the Harm Frontier as the expanding gap between what AI systems can do and what users and regulators can see, question, or control.

3 Harm Frontier

3.1 The Widening Gap

As technological progress accelerates, new forms of user privacy invasion continue to emerge. While privacy concerns in the digital age are not new, the integration of Artificial Intelligence (AI) into data collection, analysis, and inference introduces new forms of opacity and risk. Although technological developments such as privacy-enhancing technologies (PETs) and regulatory instruments (e.g. GDPR, Australia's APPs) have attempted to address privacy concerns, these responses are largely reactive, while technological innovations in AI remain proactive and fast-improving. As a result, we are witnessing a

growing misalignment between what is technologically possible and what existing legal, regulatory, and ethical frameworks can address.

3.2 Implications

The gap between what AI systems are technically capable of doing and what is legally regulated enables a growing class of data practices that are not clearly illegal yet raise serious ethical and privacy concerns. These practices operate in a grey zone where inference, profiling, and behavioural prediction occur without user awareness, meaningful consent, or oversight. Legally, however, companies are compliant with the official requirements.

This misalignment conceals the full extent of data processing within AI-integrated systems. Consumers are facing automated decision-making and continuous behavioural monitoring that extends beyond their awareness. AI systems frequently rely on passive data inputs such as browsing patterns, cursor movements, or interaction timing to generate inferences about the user. These inferences are typically not disclosed or anticipated by the user. Such invisible inference mechanisms erode individual autonomy and pose a misleading sense of control over one's data. With the lack of transparency, users cannot evaluate how their data is used, make an informed decision or resist the use of such systems. Regulatory entities, likewise, struggle to intervene effectively due to technical opacity, fragmented authority, and limited anticipatory tools.

Below, we propose a framework for analysing and visualising this growing misalignment, which we refer to as the Harm Frontier.

3.3 The Harm Frontier: Theoretical Framing

Let x be a point in time. Let $T(x)$ be the set of technological capacities at this point in time, which incorporates AI systems' ability to collect data and perform data-driven inference, and let function T represent technological change over time. Let $R(x)$ be the overall regulatory capacity to detect, prevent, or mitigate harm at the same point in time, and let function R represent regulatory change over time.

We denote $\Delta(x) = T(x) - R(x)$ as the Risk Gap, which represents the current level of privacy, ethical and social risks. The Harm Frontier emerges when the Risk Gap is positive, meaning there's a gap between technological and regulatory capabilities. When $\Delta(x) > 0$, $T(x) > R(x)$, and technologies have outpaced regulations, which puts the current state of the world in the harm zone, leading to an increased risk of harm. When $\Delta(x) = 0$, $T(x) = R(x)$, and technology and regulation are aligned, which is the desirable equilibrium. When $\Delta(x) < 0$, $T(x) < R(x)$, and regulation exceeds or anticipates technology, which likely reflects overregulation or conservative governance

3.4 Estimations

The growth of AI capabilities can be seen as dependent on computational power, algorithmic innovation, and the volume and quality of training data (Max Roser 2022). Heim et al. (2024) suggest that computing power to be the key driver for the AI process. While Moore's Law has historically described the exponential growth of computational power, Sevilla et al. (2022) showed that since 2010, this exponential growth has sped up further, to a doubling time of just about six months. Similar to previous works (Piltz et al. 2025; Heim et al. 2024), we use computational power as a proxy for capability growth, due to its quantifiable nature and strong historical correlation with system performance. Therefore, we assume that $T(x)$, reflecting the AI technological capabilities, increases exponentially.

In the meantime, regulatory changes are mostly reactive by nature and move at a much slower rate. The World Economic Forum notes that the current approach usually results in a "cat-and-mouse" game, when often corporations will look for loopholes in the law to abuse it, while others abuse legal frameworks (Malan, 2018). They note that by the time regulation is finalised, technology has typically evolved, making laws obsolete before implementation and the existing approach ineffective. The OECD has similarly warned that regulatory frameworks risk becoming obsolete due to rapid technological transformation (OECD, 2021). This misalignment between the exponential advance of technologies and an incremental move in regulation is described as the 'pacing problem' in the literature (Thierer, 2018). As Downes explains, 'technology changes exponentially, but social, economic, and legal systems change incrementally'.

Even if we assume the best possible scenario where the reaction to the new technology happens immediately, then the $R(x)$ function can be assumed as linear. Though we suggest that in real life it's logarithmic, similar to the rate of organisational change in Martec's law (see Appendix 1 for details).

The figure below shows the Harm Frontier, assuming that AI technological capacity is exponential function and regulatory capacity is linear.

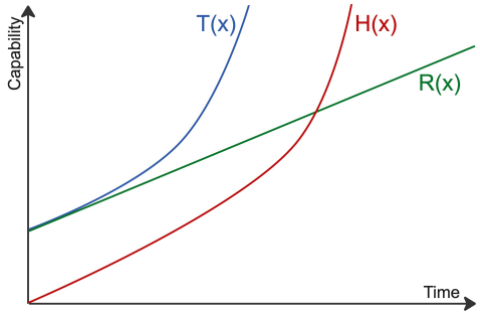


Figure 1: Harm Frontier

This function begins at zero when x = 0 and increases rapidly over time as the exponential growth in AI capability outpaces the linear trajectory of regulation. The larger H(x) gets, the greater the overseen harm potential – H(x) becomes a measure of governance failure over time. That means that the further the technological capabilities go, the greater the Harm Frontier is, unless there is a change in the regulatory processes. The shape of H(x) demonstrates how the gap between what is technologically possible and what is socially or legally governable widens at an accelerating rate.

3.5 The Harm Frontier Matrix

To visualise the Harm Frontier as a mismatch between technical capabilities for data-driven harm and the extent of governing regulation, we introduce the Harm Frontier Matrix - a two-dimensional model mapping all technology use cases into the 2-dimensional matrix. The X-axis represents the current technical capacity for harm (ranging from low to high), and the Y-axis represents current regulatory capacity (ranging from weak to strong). The Harm Frontier Matrix can be interpreted as a temporal cross-section of the broader Harm Frontier model. While the dynamic functions defined above describe how technical and regulatory capacities evolve over time, the matrix represents a snapshot at a given moment. The Harm Frontier is thus conceptualised as the gap between these two dimensions, revealing the regulatory gap in governing currently existing AI systems.

When applied to human data, AI systems are inherently positioned within a “high capacity for harm” category due to their potential to impact individual privacy, autonomy, and rights. In contrast, AI used for non-human-centred tasks such as weather forecasting or infrastructure optimisation typically carries a lower privacy harm capacity, as it operates without profiling, inference, or behavioural surveillance of individuals. Therefore, fall under the “low capacity for harm” category.

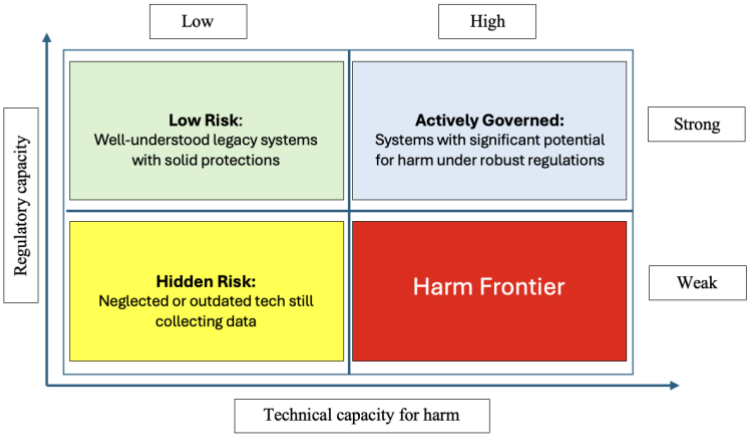


Figure 2: Harm Frontier Matrix

This matrix defines four distinct zones:

Zone 1: Low Risk. Well-understood legacy systems with limited technical capability and supported by well-established governance (e.g. static databases, printed forms, offline records).

Zone 2: Hidden Risk. Outdated or neglected systems with low technical sophistication but poor regulatory oversight, leading to unnoticed privacy risks (e.g. loyalty cards with persistent identifiers, old CRMs without updated policies).

Zone 3: Actively Governed. Digital high-capability systems with significant potential for harm, but subject to robust governance frameworks that ensure transparency, accountability, and user protection (e.g. regulated medical AI systems or AI-driven systems working with sensitive personal data).

Zone 4: Harm Frontier. Advanced systems with high technical capacity for data collection, inference, and profiling, but minimal regulatory control or transparency (e.g. AI-driven ad tech, emotion recognition tools, AI-driven user tracking). These systems exploit technical possibilities in environments where regulations lag, creating "grey zones" where users are exposed to hidden privacy risks.

By mapping digital systems across these zones, we can more clearly see where regulatory gaps are most urgent, and which types of AI-enabled systems pose the greatest latent privacy threats.

4 Empirical Evidence: Website Privacy Audit

4.1 Introduction

In this section, we illustrate the Harm Frontier by evaluating privacy risks and tracking behaviours across Australia's 50 most visited websites. This analysis sheds light on how user privacy is compromised online, often without the user's knowledge or explicit consent. By using multiple privacy auditing tools, we examine how personal and behavioural data is collected, whether through cookies, trackers, beacons, or embedded third-party services. To further illustrate how these practices interact with AI and regulatory oversight, we conduct a focused case study of a major media site (dailytelegraph.com.au).

4.2 Purpose and Scope

The data collection practices of popular websites, which we find in this project, are not unexpected; in fact, most of those practices are well-known and widespread. However, their scale combined with opaque privacy policies, integration with AI infrastructures (whether for data collection or processing) and lack of regulation around the use of AI systems within that context reflect a phenomenon described in this paper as the Harm Frontier: the growing gap between what is technically possible in data collection and inference, and what is transparently regulated. This audit helps to empirically illustrate the Harm Frontier concept by revealing how digital systems routinely exploit under-regulation of data collection that fuels or have the potential to fuel AI-driven profiling, decision-making, and surveillance, often without user consent or in ways that exploit users' lack of awareness of the AI impact on their privacy.

4.3 Methodology

In this analysis, we explore the top 50 of the most popular Australian websites[1]. The list of websites was obtained from SimilarWeb, a widely used market intelligence tool that reflects actual Australian web traffic. Then we manually classified those into 8 categories: Banking, Classified Services, Government, Job & Career, Media & Entertainment, Search & Reference, Shopping & Retail, and Social Media. For the audit of the selected websites, widely recognised privacy auditing tools were selected: Blacklight, Google Privacy Sandbox, Ghostery, EFF Privacy Badger and AESIR Privacy Scanner. The selection criteria included accessibility to non-expert users, free or low-cost usage, clarity of interface, and the transparency of results provided. All tools were intentionally chosen to reflect a consumer-facing perspective. Each tool captures several dimensions of user tracking and data collection. All tools vary in detection methods, depth, and coverage, potentially resulting in inconsistent results or partial visibility into tracking mechanisms. To mitigate these limitations, we adopted a triangulation approach, combining results across multiple scanners. This increased the reliability of tracker detection and allowed us to identify patterns that might otherwise be missed. Applying the selected tools, we were able to capture 15 distinct metrics: 1st party cookies, 3rd party cookies, blocked cookies, potential breakages, beacons, frames, trackers, ad trackers, Key logging, Session recording, Canvas fingerprinting, Facebook tracking, Google Analytics "Remarketing Audiences", Interactions with other services, and Risk Level. Each of those metrics represents the scale of data collected from a user during a website visit, and any non-zero values represent a potential violation of users' privacy. We propose using data collection

[1] The list of Australia's 50 most visited websites was curated to exclude adult content, gambling, and other sites inappropriate for general audiences. This deliberate filtering ensures the research remains broadly accessible and relevant, while maintaining alignment with general content and ethical standards.

volume as a proxy for potential AI use, based on the understanding that this aspect is one of the factors driving AI capabilities. In this study, we specifically assess the presence of third-party services integrated into websites and identify those known to possess AI capabilities. While no tools exist to reveal the presence of AI, many of the detected practices, such as extensive collection of behavioural data points, real-time personalisation and potential behavioural profiling through third-party services, are consistent with processes commonly powered by machine learning and automated inference systems. While direct observation of AI processes is also not feasible through privacy audits, the combination of high-volume data collection and AI-enabled service providers offers a reasonable basis for inferring likely AI involvement in data processing or behavioural inference.

4.4 Findings

4.4.1 Comparison of privacy auditing tools

The selected tools vary in detection methodology and scope, offering a multi-dimensional perspective on tracking behaviours across websites. A cross-tool comparison reveals substantial differences in detection capabilities, highlighting the lack of adequate tools available for end-users for effective privacy protection online and the underlying complexity of website tracking practices.

PSAT appears to be the most effective tool, identifying the highest average number of events per site and the highest cookie count. Among the tools tested, it was the most effective at revealing the scale of data collection as key evidence of the technical capability that underpins the Harm Frontier. **Blacklight** follows, particularly excelling at detecting advanced surveillance mechanisms while also adopting a user-centric approach by providing detailed descriptions of each technique and clear and accessible explanations of the results. The observed discrepancy between median and average values is likely due to some websites actively blocking the tool's scanning scripts, resulting in incomplete or failed detections. **AESIR** shows moderate tracker detection with a user-friendly risk scoring interface suitable for classifying harm frontier exposures at the site level, but it had the same issue as Blacklight, failing to scan some sites. **EFF Privacy Badger** and **Ghostery** were less effective overall, while being the only tools from our set with continuous monitoring capability as browser extensions, with Ghostery revealing the lowest average number of tracking events.

4.4.2 Key Findings

A cross-tool comparison reveals that Australian websites routinely implement high-volume tracking, which creates an opportunity for AI-driven profiling, while offering users little insight or control. These results demonstrate both the technical capability and a practical manifestation of the Harm Frontier: the sophistication of data collection technologies outpaces the visibility and governance mechanisms that are meant to protect users.

a) Average risk level. According to AESIR Privacy Scanner, 74% of Australia's top websites operate at a high privacy risk level, from a consumer perspective.

Metric results by AESIR PS	Count of websites
Total websites analysed	50
Rated "High Risk"	37
Rated "Medium Risk"	1
Rated "Low Risk"	2
Inconclusive (access blocked)	6
Inconclusive (analysis never stopped)	4

Table 1: AESIR risk level summary

b) Prevalence and Volume of Tracking. The audit revealed an extensive and pervasive presence of tracking technologies across Australia's most visited websites. PSAT, which showed to be the most effective tool in the audit, detected an average of 286 tracking-related events per site, with some websites triggering up to 665 distinct events which indicate of a deeply embedded tracking infrastructure. The extensive use of third-party cookies demonstrates the extent of cross-site tracking and data exchange. These types of cookies enable advertisers and analytics firms to

monitor user behaviour across different platforms, while collected data often contributes to AI-driven profiling and personalisation systems.

c) Frequency of fingerprinting techniques. 10 out of 38 (26%) websites analysed with Blacklight[2] had user keystroke capturing feature enabled, while 6 out of 38 (15.7%) were monitoring user mouse clicks. 8 out of 38 (21%) websites were detected to load trackers on the user's computer – a practice designed to evade third-party cookie blockers, meaning websites most likely were aware of users' attempts to preserve their privacy and still tried to prevent that.

d) Third-party interactions. Blacklight revealed a median of 8.2 third-party interactions per website, with all audited sites engaging at least one external entity and some reaching up to 21. Many of those third-party interactions involved AdTech, data brokers and tracking infrastructure providers, which indicates massive behavioural data extraction from users and shows a build-up of strong foundations for AI-driven audience segmentation and targeting. Many of these companies (e.g. LiveRamp, Neustar, Adobe) explicitly advertise services such as AI-based profiling, lookalike modelling, real-time ad targeting, data enrichment and identity resolution. Yet users are rarely informed or able to informatively consent to AI-based inference or profiling. The audit revealed that 67% (14 out of 21) of the third-party entities are linked to known regulatory issues or concerns and known for multiple privacy violations, secretive tracking and lack of transparency. A summary of this result is provided in Appendix 2.

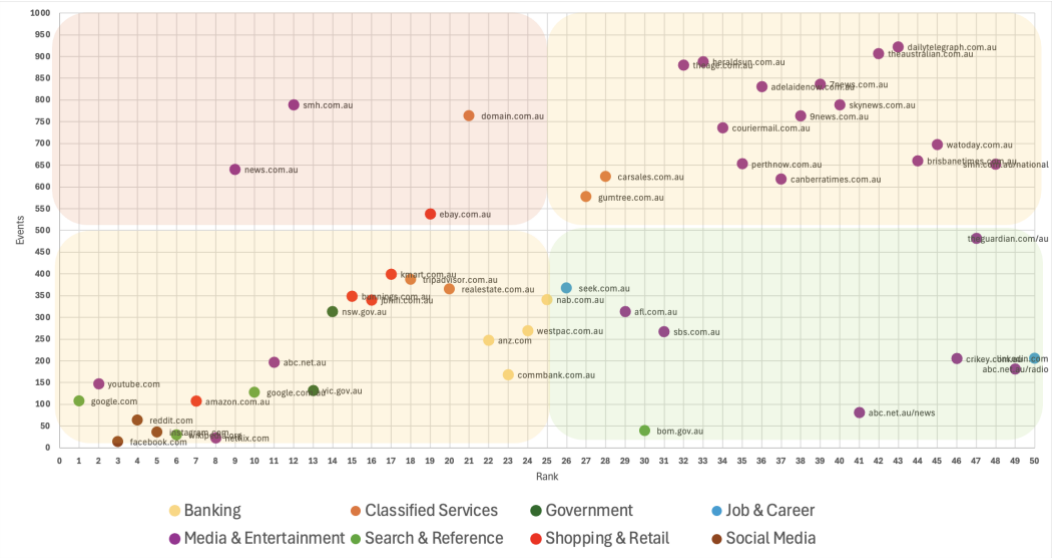
These findings emphasise the importance of the Harm Frontier – even though these websites are likely compliant with all the current regulations, the majority of the platforms possess opportunities to implement AI profiling and to invade consumer privacy.

4.4.3 Category-level findings

The category-level analysis reveals the trend that privacy risks and, therefore, the Harm Frontier are not evenly distributed across sectors. Media & Entertainment and Classified Services categories have consistently had the highest numbers of occurring "events", meaning the highest privacy risks for consumers. Social Media results are inconclusive, as most of the tools couldn't collect any events due to the requirement of logging in. Commercial platforms, particularly in media, retail, and classifieds, routinely exceed ethical boundaries, performing intensive data collection. Government and informational websites tend to operate within acceptable norms, though even those categories are not risk-free.

4.4.4 Website-level findings

The website-level analysis provides a snapshot of the current state of consumer privacy on the Top 50 websites in Australia. Each website was assigned a rank (from the Top 50 ranking list) and assessed based on the total number of privacy-relevant events detected during the audit. To contextualise and visualise the findings, each website was categorised into a privacy risk zone using a badge-based classification system.



[2] Blacklight was unable to analyse 12 websites due to access restrictions

Figure 3: Privacy risk zones

4 websites with both high rank and high occurrence of events were assigned to the Red Zone indicating a high risk for privacy. 37 fell into the Yellow Zone, representing a medium privacy risk, as they exhibited either high rank with low event count or low rank with high event count. The Green Zone included 9 websites with both low rank and low occurrence of events, indicating a lower risk for privacy. These findings suggest that 81% of the most popular Australian websites present at least moderate privacy risks to users, with only 18% operating at relatively low levels of tracking and data exposure. This aligns closely with the AESIR Privacy Scanner, which classified 74% of audited sites as "High Risk."

4.4.5 Discussion

The number and type of observed privacy events can be seen as a proxy for a website's AI-readiness. This illustrates its ability to support downstream profiling, behavioural perceptions, and manipulations. The higher the number of events, the more extensive the data collection, and the greater the AI potential for violation of user privacy. So, results from this audit allow us to measure the infrastructure that enables privacy-invasive AI. When examined through the Harm Frontier framework, many popular websites demonstrate high technical capability (T(x)) to fuel AI systems or invade user privacy, with Australian legislation (R(x)) clearly lagging behind. Although these practices are not always illegal under existing laws, they raise profound ethical and privacy concerns, meaning existing regulatory standards don't have enough capacity to detect or prevent these violations.

5 Empirical Evidence: Website Case Study

One of the websites selected for closer inspection was dailytelegraph.com.au, a major Australian news outlet operated by News Corp, which ranks among the country's most visited media platforms and in the Top-50 most popular websites in Australia. This website appeared in a Yellow Zone from the website privacy audit Matrix and had the highest number of accumulated privacy events (a summary of all measurements collected by 5 privacy auditing tools). This case illustrates a scenario where technical deployment is ambiguous for a user and is presented as useful. Still, the capabilities for behavioural surveillance with this technical implementation pose potential privacy risks. Such a discrepancy between observable practice and possible use places this case within the Red zone of the technology-regulation Matrix (Section 3.5) - the Harm Frontier, where technological potential is high and exceeds transparency and regulatory clarity, presenting the regulatory capabilities as low.

5.1 Findings from Privacy Audit Tools

Blacklight revealed 61 Ad Trackers, 45 third-party cookies, several trackers that are designed to evade third-party cookie blockers, canvas fingerprinting technique, interaction with 18 third-party services with 83% of which are known to have privacy-related regulatory violations or have previous incidents with privacy violations (refer to the table above + Microsoft (Reuters 2024), Oracle (Stuart Lauchlan 2020), Salesforce (Reuters 2024)) and at least 44% companies didn't respond or declined to answer to Blacklight on specific questions about how the company uses the data it collects from websites. PSAT recorded 594 privacy-relevant events, including 417 cookies (349 of which were third-party) from across 38 companies and 61 domains, 147 blocked cookies, 28 frames, and 2 potential breakages, placing the site among the highest tracked. Ghostery observed 25 activities and 21 potential trackers. EFF Privacy Badger observed 54 potential trackers; AESIR identified 38 cookies and 105 beacons found, assigning the site a "High Risk" score.

A closer privacy audit of dailytelegraph.com.au, a major Australian news site, revealed extensive and opaque tracking practices that exemplify the disconnect between technical capability and regulatory transparency, putting this case into the Harm Frontier. As such, the canvas fingerprinting technique was detected to be used on that website by the Blacklight tool. It is a tracking method that uses HTML5 Canvas code to generate a unique identifier for each individual user (BrowserLeaks n.d.) and relies on how different web browsers and platforms will render canvas images. While Blacklight detected a fingerprinting script being loaded on the website, it cannot determine if the purpose was user behaviour monitoring, or a less concerning activity, such as fraud prevention or bot detection. While the use of canvas fingerprinting may be justified for security-related purposes such as bot detection or fraud prevention, its presence nonetheless introduces a potential for extensive user tracking. The fact that 44% of companies declined to comment on the collected data is used and processed contributes to the lack of transparency surrounding third-party data practices. It highlights the opacity around what happens to user data once it leaves the original site.

5.2 Analysis of website Policies

5.2.1 Consent mechanism:

Empirical audit of the website use and policies audit revealed that this website uses an implied opt-out consent model – meaning the consent is implied by use. While Australian regulations allow this type of consent with some restrictions (Baig and Khan, 2023), opt-out consent is not compliant with the General Data Protection Regulation (GDPR), which generally requires opt-in consent for processing personal data (GDPR Art.7(2)). The consent mechanism on this website is an illustrative example of the Harm Frontier. Officially, the practice is compliant with regulations; however, there is a significant possibility for privacy violations as cookie consent is framed as passive, while cookies are deployed instantly. Since most users are unaware of the data collection practices, AI use to process their data and the consequences for their privacy, this is a clear case when technological capabilities pose potential privacy harms.

5.2.2 Opt-out mechanism

The opt-out interface offered by News Corp Australia raises several concerns regarding clarity, transparency, and genuine control over data collection. The opt-out mechanism is not reliably accessible across browsers and requires users to change more general browser settings that can expose users to unwanted privacy risks not only from this website, but others. The opt-out interface was not available in Safari because the "Prevent cross-site tracking" native browser feature blocked the opt-out process itself.

The toggle refers to "Personalised Ads", but not to data collection or cookie usage more broadly, creating ambiguity in consent scope. Users might interpret this opt-out as for all tracking mechanisms, when in fact third-party services may still collect and process data.

The browser compatibility check (green tick marks for JavaScript and cookies) is placed directly above the toggle. This proximity, along with technical jargon, may confuse users into associating browser readiness with consent status. The lack of separation between informational diagnostics and actionable settings reduces the interface's cognitive clarity and might be seen as a deceptive pattern.

The opt-out mechanism doesn't provide granular control for users. The interface creates an illusion of cookie control; however, in fact, users are unaware of what exact data types are included in the opt-out. The opt-out mechanism linked from the homepage directs users to a general ad preferences portal, which does not clearly explain cookie tracking or third-party data sharing. To manage cookies, users must manually navigate to a separate page, where external opt-out links (e.g., Nielsen, LiveRamp) are inconsistently placed and poorly formatted. This fragmented structure creates usability barriers, undermines meaningful consent, and places the burden on users to seek out privacy protections. Cookie opt-out practice does not meet GDPR's requirements for equal withdrawal of consent and is not aligned with OAIC's expectations for clarity and accessibility.

5.2.3 Disclosure of third parties

The privacy policy mentions some third-party entities such as Adobe, Nielsen, Ipsos, Google Analytics, and LiveRamp's Authenticated Traffic Solutions (ATS), however, analysis using Blacklight and PSAT revealed interactions with significantly more third-party trackers and data brokers, many of which are not disclosed in the policy. This lack of transparency limits users' ability to make informed privacy decisions and directly contradicts the expectations of both the OAIC's APP and international best practices. Appendix 3 provides an overview of several known regulatory issues with those third-party entities.

6 Conclusion and Future Directions

This paper introduced the Harm Frontier as a conceptual and empirical tool to capture the widening misalignment between AI-driven technological capabilities and the regulatory mechanisms designed to protect users' privacy. Through an empirical audit of Australia's top websites and a case study, we demonstrated how AI-driven tracking and profiling practices are embedded in mainstream digital services, often without user awareness or meaningful oversight. Our findings show that even legally compliant websites can contribute to significant privacy risks, particularly when AI technologies operate indirectly or remain undisclosed. The Harm Frontier framework provides a structure for analysing emerging privacy harms in digital systems in the AI era. Future work will focus on refining this framework into a practical assessment tool, integrating empirical audits with risk indicators for AI use, and identifying governance strategies to mitigate harm. The aim is to support more transparent, accountable, and privacy-aligned AI deployment in consumer environments.

7 References

ACCC. 2024. “Consumers Lack Visibility and Choice over Data Collection Practices,” , May 21.

Andreotta, A. J., Kirkham, N., and Rizzi, M. 2022. “AI, Big Data, and the Future of Consent,” AI and Society (37:4). (https://doi.org/10.1007/s00146-021-01262-5).

Ashok, M., Madan, R., Joha, A., and Sivarajah, U. 2022. “Ethical Framework for Artificial Intelligence and Digital Technologies,” International Journal of Information Management (62), p. 102433. (https://doi.org/https://doi.org/10.1016/j.ijinfomgt.2021.102433).

Battaglini, M., and Rasmussen, S. 2019. “Transparency, Automated Decision-Making Processes and Personal Profiling,” Journal of Data Protection and Privacy (2:4). (https://doi.org/10.69554/stnb6990).

BrowserLeaks. (n.d.). “Canvas Fingerprinting,” Https://Browserleaks.Com/Canvas#how-Does-It-Work.

Carroll, M., Chan, A., Ashton, H., and Krueger, D. 2023. “Characterizing Manipulation from AI Systems,” in ACM International Conference Proceeding Series. (https://doi.org/10.1145/3617694.3623226).

Chia, A. 2023. “Datafication of Work and Learning,” in Workplace Learning for Changing Social and Economic Circumstances. (https://doi.org/10.4324/9781003227946-4).

Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., Medaglia, R., Le Meunier-FitzHugh, K., Le Meunier-FitzHugh, L. C., Misra, S., Mogaji, E., Sharma, S. K., Singh, J. B., Raghavan, V., Raman, R., Rana, N. P., Samothrakis, S., Spencer, J., Tamilmani, K., Tubadji, A., Walton, P., and Williams, M. D. 2021. “Artificial Intelligence (AI): Multidisciplinary Perspectives on Emerging Challenges, Opportunities, and Agenda for Research, Practice and Policy,” International Journal of Information Management (57), p. 101994. (https://doi.org/https://doi.org/10.1016/j.ijinfomgt.2019.08.002).

Elliott, D., and Soifer, E. 2022. “AI Technologies, Privacy, and Security,” Frontiers in Artificial Intelligence (5). (https://doi.org/10.3389/frai.2022.826737).

Engström, E., and Strimling, P. 2020. “Deep Learning Diffusion by Infusion into Preexisting Technologies – Implications for Users and Society at Large,” Technology in Society (63). (https://doi.org/10.1016/j.techsoc.2020.101396).

Hadfield, G. K., and Clark, J. 2023. “Regulatory Markets: The Future of AI Governance,” ArXiv (abs/2304.04914). (https://api.semanticscholar.org/CorpusID:258060072).

Heim, L., Anderljung, M., Bluemke, E., Trager, R. 2024. “Computing Power and the Governance of AI.” GovAI, Research Post. (https://www.governance.ai/analysis/computing-power-and-the-governance-of-ai, accessed October 18, 2025).

Jurewicz, M., Klein Käfer, N., and Kran, E. (n.d.). Artificial Intelligence and Privacy: Causes for Concern.

Kashmir Hill. 2012. “How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did,” , February 16.

Kollnig, K., and Shadbolt, N. 2022. “TrackerControl: Transparency and Choice around App Tracking,” Journal of Open Source Software (7:75). (https://doi.org/10.21105/joss.04270).

Lee, H.-P. (Hank), Gao, L., Yang, S., Forlizzi, J., and Das, S. 2024. “‘I Dont Know If Were Doing Good. I Dont Know If Were Doing Bad’: Investigating How Practitioners Scope, Motivate, and Conduct Privacy Work When Developing AI Products,” in 33rd USENIX Security Symposium (USENIX Security 24), Philadelphia, PA: USENIX Association, August, pp. 4873–4890. ().

Lieber, R.2025. “Colleges Know How Much You’re Willing to Pay. Here’s How,” The New York Times. (https://www.nytimes.com/2025/05/01/business/college-tuition-price-consultants.html)

Malan, D. 2018. “The law can't keep up with new tech. Here's how to close the gap,” The World Economic Forum, (https://www.weforum.org/stories/2018/06/law-too-slow-for-new-tech-how-keep-up/)

Marks, M. 2021. “Emergent Medical Data: Health Information Inferred by Artificial Intelligence,” U.C. Irvine Law Review (11:4).

Martin, K. D., and Zimmermann, J. 2024. "Artificial Intelligence and Its Implications for Data Privacy," Current Opinion in Psychology (58), p. 101829. (https://doi.org/https://doi.org/10.1016/j.copsyc.2024.101829).

Max Roser. 2022. "The Brief History of Artificial Intelligence: The World Has Changed Fast — What Might Be Next?," Our World in Data.

Mejias, U. A., and Couldry, N. 2019. "Datafication," Internet Policy Review (8:4). (https://doi.org/10.14763/2019.4.1428).

Oluwatosin Reis, Nkechi Emmanuella Eneh, Benedicta Ehimuan, Anthony Anyanwu, Temidayo Olorunsogo, and Temitayo Oluwaseun Abrahams. 2024. "PRIVACY LAW CHALLENGES IN THE DIGITAL AGE: A GLOBAL REVIEW OF LEGISLATION AND ENFORCEMENT," International Journal of Applied Research in Social Sciences (6:1). (https://doi.org/10.51594/ijarss.v6i1.733).

Panoptykon Foundation. 2021. "Algorithms of Trauma: New Case Study Shows That Facebook Doesn't Give Users Real Control over Disturbing Surveillance Ads," , October 6.

Reece, A. G., and Danforth, C. M. 2017. "Instagram Photos Reveal Predictive Markers of Depression," EPJ Data Science (6:1). (https://doi.org/10.1140/epjds/s13688-017-0110-z).

Reece, A. G., Reagan, A. J., Lix, K. L. M., Dodds, P. S., Danforth, C. M., and Langer, E. J. 2017. "Forecasting the Onset and Course of Mental Illness with Twitter Data," Scientific Reports (7:1). (https://doi.org/10.1038/s41598-017-12961-9).

Reuters. 2024. "EU Privacy Regulator Fines LinkedIn 310 Mln Euro," , October 25.

Romansky, R. P., and Noninska, I. S. 2020. "Challenges of the Digital Age for Privacy and Personal Data Protection," Mathematical Biosciences and Engineering (17:5), pp. 5288–5303. (https://doi.org/10.3934/mbe.2020286).

Russell, S., Norvig, P. (2021). Artificial Intelligence: A Modern Approach. Pearson. (http://lib.ysu.am/disciplines_bk/efdd4d1d4c2087fe1cbe03d9ced67f34.pdf, accessed October 18, 2025).

Sartor, G., and Lagioia, F. 2020. "The Impact of the General Data Protection Regulation (GDPR) on Artificial Intelligence," European Union.

Sevilla, J., Heim, L., Ho, A., Besiroglu, T., Hobbhahn, M. and Villalobos, P. 2022. "Compute Trends Across Three Eras of Machine Learning," 2022 International Joint Conference on Neural Networks (IJCNN), Padua, Italy, pp. 1-8. (doi: 10.1109/IJCNN55064.2022.9891914).

Sheikh, H., Prins, C., Schrijvers, E. 2023. "Artificial Intelligence: Definition and Background. In: Mission AI. Research for Policy." Springer, Cham. (https://doi.org/10.1007/978-3-031-21448-6_2).

Solow-Niederman, A. 2022. "INFORMATION PRIVACY AND THE INFERENCE ECONOMY," Northwestern University Law Review (117:2). (https://doi.org/10.2139/ssrn.3921003).

Stuart L. 2020. "Oracle and Salesforce Hit by €10 Billion Privacy Violation Class Action - and There's an Elliott Management Angle to Stir the Pot Further," Diginomica.Com , , August 14.

Thierer, A. 2018. "The Pacing Problem and the Future of Technology Regulation," Mercatus Center, (https://www.mercatus.org/economic-insights/expert-commentary/pacing-problem-and-future-technology-regulation)

TikTok Privacy Center. (n.d.). "Ads and Your Data."

Appendix 1 – Martec's Law

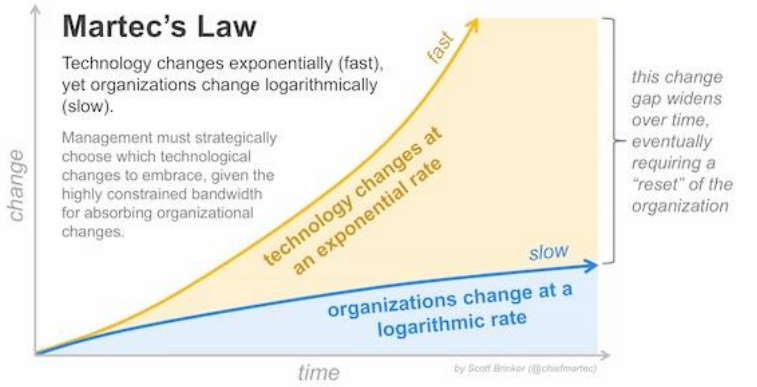


Figure A.1: Martec's Law

Appendix 2 – Website Privacy Audit protocol

To ensure consistent and replicable testing conditions, all website audits were conducted using the Google Chrome browser (version 134) and Safari browser (version 18.3) on macOS (version 15 Sequoia) with all built-in privacy protections manually disabled to allow complete visibility into tracking behaviours.

Google Chrome

1. Prevent cross-site tracking (off)
2. Block all cookies (off)
3. Use advanced tracking and fingerprint protection (off)
4. Hide IP address from trackers (off)

Safari:

1. Enhanced protection (off, set to Standard protection)
2. Allow third-party cookies (on)
3. Send a 'Do Not Track' request with your browsing traffic (off)

Each website was accessed in private browsing (incognito) mode to avoid interference from stored cookies, cached data, or prior browsing history.

For each website, a suite of privacy auditing tools was run in a way that all 50 websites were consequently evaluated with each tool in the following order:

1. Blacklight (Safari)
2. Google Privacy Sandbox (Chrome)
3. Ghostery (Safari)
4. EFF Privacy Badger (Safari)
5. AESIR Privacy Scanner (Safari)

Only one tool was applied (or enabled) at a time to avoid interference. In cases where a tool failed to return results (due to transient issues or internal processing errors), the audit was repeated to obtain valid output. Some tools failed to provide results for some of the websites due to the login requirements or additional protections employed by websites, such as bot or crawler protections.

All findings were recorded immediately after the tool executed, and results were interpreted using each tool's published guidance and documentation. The methodology for each tool can be found in the following sources:

1. Blacklight, https://themarkup.org/blacklight/2020/09/22/how-we-built-a-real-time-privacy-inspector
2. Google Privacy Sandbox, https://privacysandbox.google.com
3. Ghostery, https://www.ghostery.com/blog/how-we-at-ghostery-protect-users-from-web-tracking#protecting-our-users-from-web-tracking
4. EFF Privacy Badger, https://privacybadger.org
5. AESIR Privacy Scanner, https://aesirx.io/blog/privacy-scanner/what-is-aesirx-privacy-scanner-and-how-does-it-work

All experiments were conducted between the 21st and 28th of March 2025.

Appendix 3 – Website Privacy Audit results

Rating	Tools	Tracking mechanism	Median	Average	Max
1	PSAT	1st cookies	36	38	88
		3rd cookies	108	161.3	412
		Blocked cookies	79	77.2	169
		Frames	7.5	13.1	34
		Breakages	0	0.6	11
		Sum of events	**220.5**	**286.1**	**665**
2	Blacklight	Add trackers	14	26.2	72
		3rd p cookies	21.5	50.6	223
		Sum of events	**27**	**74.1**	**295**
3	AESIR	3rd cookies	21.5	20.9	52
		Beacons	27.5	38.9	105
		Risk level			
		Sum of events	**21.5**	**47.9**	**144**
4	EFF	Potential Trackers	16	19.9	55
		Sum of events	**16**	**19.9**	**55**
5	Ghostery	Trackers	8	8.9	29
		Observed activities	10	11.4	31
		Sum of events	**10**	**11.2**	**31**
		Aggregate number of events	**357**	**421.4**	**922**

Table A.1: Website Privacy Audit results

Appendix 4 – Regulatory issues of News Corp Australia third-party partners

Name	Function	Known regulatory issues
Adobe	Adtech and analytics platform; provides audience measurement and targeting	Subject of concerns over tracking in Adobe Experience Cloud and suit filled in Dutch (2025) alleging it built profiles of internet users through secretive tracking, and profited by selling the data (McCroskey 2025; Zwets 2023).
Alphabet	Parent company of Google; provides search, ad services, analytics (e.g., DoubleClick, Google Ads)	Multiple GDPR complaints (BigId n.d.; Reuters 2025). Fined €57M in France (EDPB, 2019) for multiple GDPR violations, including a lack of transparency and valid user consent for personalized advertising data processing
Amazon	Operates ad services (Amazon Ads), demand-side platform, and cloud analytics	Investigated for GDPR violations; fined €746M in Luxembourg for failing to obtain explicit consent for the use of customer data in targeted advertising (Syrenis 2025).
Comscore	Audience measurement and media analytics	Settled SEC charges for fraud (Debevoise & Plimpton LLP 2019) - charges focused on financial misrepresentation rather than direct privacy violations; general privacy concerns; class action against the company accusing online data tracking service of collecting confidential information, unauthorised scanning of files on users' personal computers and modifying their security settings. Class action lawsuit filed (Mirando 2011).
Criteo	Behavioural retargeting and real-time bidding for ads	Fined €40 million by the French data protection authority (Mikalef 2023) for GDPR violations including failing to obtain valid consent for tracking, lacking transparency in its data practices, and failing to adequately provide users with the right to access and erase their personal data
DoubleVerify	Ad verification and fraud detection service	No known issues
IPONWEB	Programmatic ad infrastructure and AI-based bidding platform	Acquired by Criteo; subject to same regulatory ecosystem scrutiny
LinkedIn	Social media platform; provides behavioural ad targeting through Microsoft Ads	Fined €310M in Ireland (Reuters 2024) for behavioral advertising and targeted ads without obtaining freely given, informed, and unambiguous consent. Known to enable use of users' data for AI training by default with opt-out consent model (Arntz 2025).
LiveRamp	Identity resolution and cross-device user tracking	Faces privacy concerns (TaulerSmith n.d.; ORG 2024) due to its business model, which involves collecting and processing vast amounts of personal and pseudonymous data to facilitate identity resolution and targeted advertising across various online platforms. Currently facing lawsuit, Accused of Collecting and Selling Consumer Data to Third Parties Without Consent
Lotame	Data Management Platform (DMP); creates behavioural profiles and audience segments	Flagged in privacy watchdog reports

MediaMath	Demand-side platform for programmatic advertising	No known issues
Neustar	Identity resolution and marketing analytics (acquired by TransUnion)	Acquired by TransUnion. Faced accusations of privacy violations (DataGuidance 2022) when Senator Ron Wyden requested a Federal Trade Commission (FTC) investigation into its Security Services division. Accused of allegedly selling sensitive metadata of U.S. citizens to a Department of Defense project at Georgia Tech, potentially violating promises made to consumers and the FTC Act of 1914.
OpenX	Programmatic advertising exchange	Fined $2M by FTC for collecting data from children (COPPA violation) (FTC 2021).
PubMatic	Supply-side platform; real-time bidding and ad optimisation	No known issues
Quantcast	AI-powered audience measurement and targeting	Subject to ICO and NGO scrutiny; Subject to complaints by the advocacy group for the way they gathered, used, and sold users' personal data for ad revenue without their knowledge or consent (Xcina n.d.).
Magnite	Supply-side platform for publishers; ad monetisation	Formerly Rubicon Project; No known issues
Sovrn	Ad monetisation and affiliate marketing platform	No known issues
TowerData	Email intelligence and identity matching (data broker functions)	No known issues
Verizon	Operates Verizon Media (now Yahoo); provides adtech and behavioural data services	Past fines and scrutiny for illegally sharing access to consumers' location data without obtaining their consent and also failed to take "reasonable measures" to safeguard customers' location information (Nya 2024; CNBC 2025).
WPP	Advertising and public relations holding company; owns data marketing subsidiaries	No known issues
Zeta Global	AI-powered marketing platform; personalisation and customer lifecycle management	Flagged for opaque practices in data enrichment and profiling, under investigation by the U.S. (KTMC 2024), class action lawsuit filed for "predatory consent farms" and misleading financial statements due to artificially inflated data collection practices (Labaton 2025).

Table A.2: Regulatory issues of News Corp Australia third-party partners

References for Appendix 4:

Arntz, P. 2025. "LinkedIn will use your data to train its AI unless you opt out now." (https://www.malwarebytes.com/blog/news/2025/09/linkedin-will-use-your-data-to-train-its-ai-unless-you-opt-out-now, accessed October 18, 2025)

BigId. N.d. "TheStreet: Alphabet's $57 Million Data Fine: Drop in the Bucket, or Sign of Trouble Ahead?" (https://bigid.com/news/thestreet-alphabets-57-million-data-fine-drop-in-the-bucket-or-sign-of-trouble-ahead/, accessed October 18, 2025)

Burgess, M. 2021. "Amazon's Massive GDPR Fine Shows the Law's Power—and Limits." (https://www.wired.com/story/amazon-gdpr-fine-shows-power-limits/, accessed October 18, 2025)

CNBC. 2025. "U.S. court upholds Verizon $46.9 million fine over location data." (https://www.cnbc.com/2025/09/10/us-court-upholds-verizon-46point9-million-fine-over-location-data.html, accessed October 18, 2025)

DataGuidance. 2022. "USA: Senator requests FTC investigation into Neustar Security Services sale of metadata." (https://www.dataguidance.com/news/usa-senator-requests-ftc-investigation-neustar-security, accessed October 18, 2025)

Debevoise & Plimpton LLP. 2019. "Media Analytics Firm and Former CEO Settle Revenue Recognition Charges." (https://www.debevoise.com/insights/publications/2019/10/media-analytics-firm-and-former#:~:text=On%20September%2024%2C%202019%2C%20Comscore,February%202014%20through%20February%202016, accessed October 18, 2025)

EDPB. 2019. "The CNIL's restricted committee imposes a financial penalty of 50 Million euros against GOOGLE LLC." (https://www.edpb.europa.eu/news/national-news/2019/cnils-restricted-committee-imposes-financial-penalty-50-million-euros_en#:~:text=On%2021%20January%202019%2C%20the,consent%20regarding%20the%20ads%20personalization., accessed October 18, 2025)

Epiq. n,d. "Recent GDPR Fines Against Amazon and WhatsApp Set New Records." (https://www.epiqglobal.com/en-gb/resource-center/articles/recent-gdpr-fines-against-amazon-and-whatsapp, accessed October 18, 2025)

FTC. 2021. "Advertising Platform OpenX Will Pay $2 Million for Collecting Personal Information from Children in Violation of Children's Privacy Law." (https://www.ftc.gov/news-events/news/press-releases/2021/12/advertising-platform-openx-will-pay-2-million-collecting-personal-information-children-violation, accessed October 18, 2025)

KTMC. 2024. "ZETA GLOBAL HOLDINGS CORP. (NYSE: ZETA) SECURITIES FRAUD CLASS ACTION." (https://www.ktmc.com/new-cases/zeta-global-holdings-corp#:~:text=This%20is%20a%20securities%20fraud%20class%20action,purchased%20or%20acquired%20Zeta%20Global%20Holdings%20Corp, accessed October 18, 2025)

Labaton. 2025. "Davoodi v. Zeta Global Holdings Corp., No. 24-8961 (DEH) (S.D.N.Y.). "(https://www.labaton.com/cases/davoodi-v-zeta-global-holdings-corp, accessed October 18, 2025)

McCroskey, K. 2025. "Class Action Lawsuit Claims Adobe Secretly Intercepts, Monetizes Consumer Data via Online Tracking Tech." (https://www.classaction.org/news/class-action-lawsuit-claims-adobe-secretly-intercepts-monetizes-consumer-data-via-online-tracking-tech, accessed October 18, 2025)

Mikalef, W. 2023. "CNIL fines adtech company Criteo €40 million – what are the key takeaways?" Bird&Bird, (https://www.twobirds.com/en/insights/2023/global/cnil-fines-adtech-company-criteo-40-million-what-are-the-key-takeaways, accessed October 18, 2025)

Mirando, S. 2011. "comScore Privacy Class Action Lawsuit." (https://topclassactions.com/lawsuit-settlements/lawsuit-news/comscore-privacy-class-action-lawsuit/, accessed October 18, 2025)

Nya, A. 2024. "AT&T, Sprint, T-Mobile, and Verizon Fined for Failing to Protect Our Location Data." Goof Jobs First. (https://goodjobsfirst.org/att-sprint-t-mobile-and-verizon-fined-for-failing-to-protect-our-location-data/, accessed October 18, 2025)

ORG. 2024. "ORG SUBMITS COMPLAINTS ABOUT INTRUSIVE LIVERAMP ADTECH SYSTEM." (https://www.openrightsgroup.org/press-releases/org-complaint-liveramp-adtech/, accessed October 18, 2025)

Reuters. 2024. "EU privacy regulator fines LinkedIn 310 mln euro." (https://www.reuters.com/technology/eu-privacy-regulator-fines-linkedin-310-mln-euro-2024-10-24/, accessed October 18, 2025)

Reuters. 2025. "Alphabet hit with Austrian privacy complaint over alleged browser tracking." (https://www.reuters.com/technology/alphabet-hit-with-austrian-privacy-complaint-over-alleged-browser-tracking-2024-06-13/ , accessed October 18, 2025)

Syrenis. 2025. "Luxembourg upholds €746m Amazon fine." (https://syrenis.com/resources/blog/luxembourg-amazon-fine/, accessed October 18, 2025)

TaulerSmith. n.d. "Invasion of Privacy Lawsuit Against LiveRamp." (https://www.taulersmith.com/invasion-of-privacy-lawsuit-against-liveramp/, accessed October 18, 2025)

Xcina. n.d. "Second wave of complaints targets deceptive cookie banners." (https://xcinaconsulting.com/blog/second-wave-of-complaints-targets-deceptive-cookie-banners/, accessed October 18, 2025)

Zwets, B. 2023. "Adobe sued in mass claim for illegal tracking." (https://www.techzine.eu/news/privacy-compliance/114296/adobe-sued-in-mass-claim-for-illegal-tracking/, accessed October 18, 2025)

Copyright

Copyright © 2025 Nataliia Tashevtseva, Dr Suelette Dreyfus, Prof Udaya Parampalli. This is an open-access article licensed under a Creative Commons Attribution-Non-Commercial 4.0 International License, which permits non-commercial use, distribution, and reproduction in any medium, provided the original author and ACIS are credited.