A Year Under the DSA: Ad Transparency's Uneven Landscape

Abir Benzaamia
LIX, CNRS, Inria, École Polytechnique, IP Paris
Palaiseau, France
abir.benzaamia@inria.fr

Asmaa El fraihi
LIX, CNRS, Inria, École Polytechnique, IP Paris
Palaiseau, France
asmaa.el-fraihi@inria.fr

Ines Abdelaziz
LIX, CNRS, Inria, École Polytechnique, IP Paris
Palaiseau, France
ines.abdelaziz@inria.fr

Oana Goga
LIX, CNRS, Inria, École Polytechnique, IP Paris
Palaiseau, France
oana.goga@inria.fr

Abstract

The Digital Services Act (DSA) has put platform accountability on center stage, requiring online platforms to provide greater transparency into how advertisements are targeted and delivered to users. Central to these obligations are two mechanisms: **user-facing ad explanations**, which inform individuals why they were shown a given ad, and **public ad repositories**, which are intended to enable independent auditing of advertising practices. This study provides the first multi-platform evaluation of these two mechanisms across Facebook, Instagram, YouTube and X. Using 48,511 user-facing *"Why am I seeing this ad?"* (WAIST) notices, and a systematic analysis of each platform's public ad repository, we assess how well current implementations disclose the parameters and decision processes involved in targeting. To do so, we develop and apply an operational framework based on Articles 26 and 39 of the DSA—capturing the granularity, attribution of targeting and delivery choices, data source disclosures, and accuracy—and apply it across both user-facing notices and public ad repositories. Our findings show that transparency remains fragmented and inconsistent across platforms. User-facing explanations vary widely in precision and often omit key targeting information, while repositories provide incomplete, misattributed, and at times difficult-to-interpret targeting data. Moreover, discrepancies between explanations and repository entries undermine the reliability of both mechanisms. Overall, current transparency infrastructures fall short of the DSA's expectations and highlight the need for clearer and more enforceable standards for advertising transparency moving forward.

Keywords

Digital Services Act (DSA), Online Advertising, Ad transparency, Ad explanations, Ad repositories.

1 Introduction

Online advertising has become one of the most regulated components of the digital ecosystem. Growing public distrust over how platforms collect, use, and monetize personal data through advertising has prompted policymakers to push for stronger rules governing the transparency and accountability. In the European Union (EU), the *Digital Services Act (DSA)* [15]—introduced in 2022 and fully applicable since February 2024—sets out new requirements for online platforms aimed at creating a safer and more transparent environment for both users and society at large.

These requirements are particularly relevant given how online advertising systems actually operate. Platforms such as Meta, Google, and X rely on complex targeting systems designed to maximize relevance for both advertisers and users. Targeting draws on a wide range of signals, from contextual cues such as time, location, and the content being viewed, as well as user profiles derived from declared demographics, behavioral patterns within and beyond the platform, and inferred interests.

To meet transparency obligations, the DSA requires platforms to provide two distinct forms of disclosure. The first are *ad explanations* (Article 26) [13], presented directly to users, which are intended to clarify why a particular ad was shown and what data informed that decision. In principle, explanations should be clear, interpretable, and actionable, enabling users to understand targeting decisions and, where desired, to adjust them through interfaces such as *"Why am I seeing this ad?"* notices and ad preference settings. These explanations should also reveal the advertiser's intended audience, and the platform's reasoning, including how data sources (e.g., browsing activity, in-platform behavior, search history) and inferences were applied. The second mechanism is the provision of *public ad repositories* (Article 39) [14], or ad libraries, which are intended to support scrutiny beyond the individual user. These repositories are meant to provide searchable records of ads that have run on the platform, including information on targeting criteria, advertiser identity, and delivery parameters.

While explanations are intended to give more transparency and control for users, ad repositories are designed to enable accountability by regulators, journalists, and researchers. Together, these two mechanisms should provide both personalized insights for users and system-level transparency for external auditors.

Although the DSA mandates unprecedented levels of transparency, it remains unclear whether the current platform implementation translates into meaningful accountability. Prior research preceding the DSA indicates that ad explanations can be vague, incomplete, or misleading [3, 40, 41], and that public repositories may lack the detail needed to detect problematic ads [12]. At the

Proceedings on Privacy Enhancing Technologies 2026(2), 517–532

https://doi.org/10.56553/popets-2026-0059

same time, platforms continue to rely on opaque optimization systems that dynamically adjust targeting [1, 27]. These dynamics raise the question: *To what extent do current ad transparency mechanisms across major platforms provide meaningful insight for users and auditors into how online advertisements are targeted and delivered?*

To investigate this, we assess how the two core ad transparency mechanisms introduced by the regulation are implemented across major platforms through a case study on **Germany**, by analyzing data collected through a browser extension that captures ads and their associated targeting explanations found in *"Why Am I seeing this ad?"* (WAIST) notices, as experienced by real users (Sections 7 and 8). Our analysis covers $N = 48,511$ ad explanations collected across four platforms: **Facebook**, **Instagram**, **YouTube**, and **X (formerly Twitter)**. In addition, we examine the platforms' public ad repositories and compare the targeting details shown to users with those made public (Section 10). We also assess compliance with Articles 26 [13] and 39 [14] of the DSA and derive operationalized criteria from these provisions to clarify how advertising transparency should be implemented moving forward.

Our findings show that advertising transparency remains uneven across platforms, with gaps appearing in both user-facing explanations and public ad repositories. First, **user-facing explanations differ in precision and reliability**. Instagram and Facebook offer relatively consistent and highly granular notices, whereas YouTube provides largely generic statements that vary across profiles. X falls between these two extremes but still exhibits some inconsistencies (e.g., *"The reason why you were shown this ad is not available"*). Second, **platforms rarely provide clear disclosure of data sources or responsibility for targeting decisions**. Although Instagram and Facebook include UI elements intended to distinguish advertiser choices from platform-generated inferences, this distinction is mostly made for behavioural targeting (interest-based advertising). On YouTube and X, however, this distinction is not made at all. All platforms only vaguely reference data sources, and do not provide actionable details that would enable users to exercise control over their data. Controlled experiments on YouTube further show that certain targeting criteria are omitted from WAIST notices altogether (Section 8). Third, **public ad repositories provide incomplete and often misleading accounts of how ads are targeted**. Meta discloses only demographic and location attributes; Google lists broad categories that do not correspond to the full range of advertiser options; and X exposes detailed but unstructured targeting tokens that lack labeling and are difficult to interpret. Across all three platforms, repositories attribute all targeting parameters to advertisers—even when controlled experiments demonstrate that some parameters were introduced by platform optimization systems (Section 10.3). Moreover, repository data frequently fails to align with user-facing explanations. On Google and X, substantial portions of the targeting attributes shown to users are missing from repository entries (46.5% and 28.9%, respectively).

Together, these inconsistencies obscure how ads are actually targeted and limit the effectiveness of current implementations of the DSA's ad transparency provisions. They also highlight the need for clearer and actionable standards for advertising transparency.

2 Background

This section presents an overview of online advertising and how its ecosystem is regulated by the DSA in terms of transparency.

2.1 Online Advertising

Online advertising operates through a sophisticated and complex ecosystem of *advertisers*, *ad platforms*, and *algorithmic decision-making*. Large platforms like Google and Meta, track users' browsing activity across the web and create user profiles for the purpose of sending them targeted advertising. These online platforms have rich data for developing extensive user profiles [3, 4, 39], augmenting website visits with user-provided personal information and interactions with platform content (likes, clicks, video views).

Advertisers rely on these ad platforms to reach their audiences, where they explicitly select their **targeting parameters** including targeted audience criteria (demographics, interests, behaviors, etc), as well as their objectives (e.g., maximizing clicks, conversions, ROI, etc.) [27]. Platforms implement Real Time Bidding (RTB) auctions to sell impressions to advertisers [5, 11], and determine **the delivery** of their ads—deciding which ads to show to which users and when. This algorithmic decision-making combines contextual factors (time of day, content being viewed, device type, placement position) with inferred user profiles. The result is a prediction-driven system where algorithms continuously adjust **targeting** and **delivery** strategies to maximize engagement and advertising revenue [8, 19, 20, 27, 30]. This complexity introduces inherent opacity. A given ad shown to a user may reflect *advertiser-defined targeting parameters*, *platform-side delivery optimizations*, or both.

Recently, the selection of these criteria has become increasingly delegated to AI systems that autonomously determine which attributes, interests, and behaviors are most predictive of engagement [32, 37]. While AI has always played a role in *delivery optimization*, this transition in particular represents a structural shift away from advertiser-led targeting choices toward full *algorithmic-driven targeting* [9, 10].

2.2 The Digital Services Act (DSA)

The DSA is an EU regulation that sets binding rules for advertising transparency. Articles 26 and 39 [13, 14] specify obligations for online platforms, including Very Large Online Platforms (VLOPs)—platforms with a user base of over 45 million in the (EU)—that must meet heightened standards. Two mechanisms are essential to advertising transparency: **user-facing ad explanations** and **public ad repositories**.

User-facing Ad Explanations. Article 26 requires online platforms to provide users with *"meaningful information"* about why a particular advertisement is shown[1]. This includes: the identity of the

[1]DSA Article 26(1) (bolded by the authors) "Providers of online platforms that present advertisements on their online interfaces shall ensure that, for each specific advertisement presented to each individual recipient, the recipients of the service are able to identify, in a clear, concise and unambiguous manner and in real time the following: (d) **meaningful information** directly and easily accessible from the advertisement about **the main parameters** used to determine the recipient to whom the advertisement is presented" [13]

advertiser, the main targeting parameters used, and whether third-party data or profiling was involved. These explanations are typically presented to users through the *"Why am I seeing this ad?"* (WAIST) information notice that accompanies ads [21, 33, 34, 42].

Public Advertising Repositories. Article 39 requires platforms designated as Very Large Online Platforms (VLOPs) to maintain public repositories (also called ad libraries or transparency centers) of all ads shown on their services for at least a duration of a year after they stop running[2][3]. These repositories are meant to support accountability and auditing by making ads searchable by content, advertiser, targeting criteria, and delivery metrics.

2.3 "Why am I seeing this ad?" (WAIST)

The "Why am I seeing this ad?" (WAIST) notice is a user-facing transparency feature first introduced by Meta (formerly Facebook) in 2014 [6]. It was designed to provide users with a brief explanation of why a particular advertisement appears in their feed—typically referencing factors such as demographics, inferred interests, or prior interactions with the advertiser. Following Meta's introduction of this feature, similar notices were gradually adopted by other major platforms, including Google (YouTube) and X (formerly Twitter), as public concerns grew regarding opaque advertising, personalization, and targeting practices.

In their current form [21, 33, 34, 42], WAIST notices generally include: (1) **information about the advertiser**; (2) the main **targeting rationale** explaining why the user was selected to see the ad; and (3) **user controls**, such as options to hide the ad, view or edit interest topics, or manage personalization preferences. Examples of these notices are shown in Figures 3, 4, and 5 in the Appendix.

In this work, we focus exclusively on the **targeting rationale** component, namely, the text explanations (or phrases) that describe the parameters chosen by the advertiser or the signals inferred by the platform to target a given user with an ad.

3 Related Work

Growing public scrutiny and regulatory pressure have pushed online advertising platforms to introduce a variety of transparency mechanisms, including user-facing *ad explanations*, *ad preferences pages*, and *public ad repositories* [3, 4, 9, 17, 18]. Researchers and regulators have examined whether these tools actually enable users and auditors to understand how ads are targeted and delivered. Early studies have revealed a persistent gap between the targeting information shown to users and the targeting capabilities available to advertisers [2, 3, 40, 41]. Andreou et al. [3] found that Facebook's WAIST notices were often incomplete or misleading, failing to disclose the full set of attributes used in targeting. Subsequent research showed that advertisers routinely employ a wide range of parameters—including sensitive attributes, partner-provided data, and offline sources—that are rarely reflected in these explanations [2, 40, 41]. This body of work has informed subsequent policy interventions, culminating in the regulatory standards now codified in the Article 26 of the DSA [13].

Beyond ad explanations, platforms such as Google and Meta offer *ad preferences* managers (e.g., Google Ad Center [24] and Facebook Ad Preferences [16]) that display interest categories inferred about users. However, these explanations are typically vague and inconsistent, omitting details on data provenance or inference logic, and providing only a superficial picture of user profiling and micro-targeting [4, 38, 41].

Similar limitations are found in platform-level ad repositories, which are intended to support public oversight. In terms of political advertising transparency, Edelson et al. [12] documented significant shortcomings in Meta's Ad Library, showing that advertisers could easily manipulate disclosures or evade classification as political ads. The Mozilla Foundation's 2023 *Ad Transparency Stress Test* [17] further demonstrated that repositories across twelve major platforms suffer from limited accessibility, missing data, and inconsistent functionality. Using criteria inspired by the DSA Article 39 [14] and expert guidelines [35], Mozilla concluded that no major platform provides a repository that is adequate for systemic auditing or regulatory monitoring.

Our study extends on this body of work by returning to the question of what transparency should concretely entail. Rather than proposing new normative standards, we systematically benchmark platform transparency against the DSA's requirements. We operationalize our criteria directly from the regulation and evaluate how current implementations of *user-facing* ad explanations and *public* ad repositories convey information about targeting parameters selected by advertisers or inferences made by platforms using user data. In doing so, our analysis identifies where the transparency tools of VLOPs align or fall short of the DSA's vision of meaningful accountability.

4 Dataset

We rely on a dataset of $N = 48,511$ ads collected through *Who Targets Me (WTM)* [28], a browser extension that records ads shown to users along with their associated *"Why am I seeing this ad?"* (WAIST) explanations (see Appendix B for details). The collection covers four major platforms: Facebook, Instagram, YouTube, and X (formerly Twitter), with all data gathered from users based in **Germany**, where the DSA applies. For Facebook, data spans from **September 2017** to **March 2025**, though with occasional inconsistencies and gaps due to interface changes that intermittently disrupted automated collection. For Instagram, YouTube, and X, data was collected between **September 2024** and **March 2025**.

Sampling Strategy. From the WTM corpus, we constructed a random sample for analysis. Specifically, for each platform, we sampled 500 ads per month from September 2024 to March 2025, among active users (defined as users who viewed at least 20 ads). This sampling yields **3,500 WAIST notices per platform**—for Instagram, YouTube, X, and Facebook—which we use to benchmark the transparency landscape against the DSA in terms of both user-facing explanations and ad repositories (see Sections 7, 8, and 10).

[2]DSA Article 39(1) "Providers of very large online platforms or of very large online search engines that present advertisements on their online interfaces shall compile and make publicly available[..]a repository containing the information referred to in paragraph 2" [14]

[3]DSA Article 39(2) "The repository shall include at least all of the following information[..] (e) whether the advertisement was intended to be presented specifically to one or more particular groups of recipients of the service and if so, **the main parameters** used for that purpose including where applicable the main parameters used to exclude one or more of such particular groups." [14]

An additional dataset was constructed with all available Facebook data, comprising **38,011 WAIST notices** from September 2017 to March 2025. We use this dataset for conducting a longitudinal analysis on Facebook (see Section 9).

Data Structure. The format of WAIST data acquired differs by platform. On Facebook, data was stored as structured JSON objects, collected by intercepting network calls when an ad is displayed (see Appendix A.1 for an example). On Instagram, YouTube, and X, data is captured as raw HTML. We preprocess all datasets to extract the **text targeting rationale** components, which serve as the basis for our analysis.

4.1 Categorization of Targeting Parameters

We classify the text targeting explanations found in WAIST notices into targeting categories. This is necessary to standardize heterogeneous explanation formats and compare targeting and transparency practices across platforms.

4.1.1 Methodology. To define the targeting categories, we reviewed prior research on online advertising and targeting systems [29]. We also examined official platform documentation and advertiser help pages [22, 44] to determine which targeting options are available and how they are described. Finally, we inspected *ad manager* interfaces directly on each platform, by creating advertiser accounts on Meta [31], Google [25] and X [43], and verifying where specific attributes appear in the advertiser workflow. Based on this analysis, we defined seven targeting categories: **Contextual**, **Demographic**, **Time and Location**, **Behavioral**, **Retargeting**, **Custom Audience**, and **Lookalike**. Each targeting explanation text extracted from the WAIST notices was assigned to one of these categories using a **rule-based text classification** procedure. We constructed regular expression patterns to detect characteristic phrases in the explanation text that indicate the targeting mechanism. For example, phrases that mention *"interests"* or *"past activity"* were mapped to behavioral targeting, or *"shown to people living in"* to location-based targeting. Table 1 lists targeting categories and their corresponding descriptors and example patterns.

4.1.2 Prevalence of Targeting Categories. Table 2 presents the distribution of text explanations within each targeting category. The most prevalent categories across platforms in our dataset are demographic and time and location targeting. References to contextual targeting cues are rare, appearing only occasionally on YouTube (1.12%) and even less frequently on X (0.28%). We also find that behavioral (interest-based) targeting is referenced most often on Instagram (33.22%), followed by YouTube (12.66%), X (9.31%), and Facebook (7.86%)[4]. Mentions of retargeting, custom audiences, and lookalike audiences occur sporadically on Facebook, Instagram, and YouTube, but are substantially more common on X. These results highlight clear differences in how platforms represent and disclose targeting parameters, reflecting both the underlying advertising systems and their respective transparency implementations.

[4]Due to partial scraping coverage, Facebook's values should be interpreted with caution.

5 Ethics

Data were collected exclusively from consenting users who voluntarily installed the Who Targets Me (WTM) browser extension, which collects ads and their explanations. Our data sample **did not include any personally identifiable information** (PII), and our analysis focuses on the content of ad explanations and *public* ad archives. As such, our work does not involve interventions with human subjects and does not require institutional review board (IRB) approval. This work is, however, covered by a project-level IRB approval at our institution.

6 Operationalizing DSA-mandated Transparency

To benchmark the current advertising transparency landscape, we operationalize the criteria in Articles 26 and 39 of the DSA involving disclosure of targeting details. We then use these operational definitions to evaluate ad explanations and ad repositories in Sections 7 and 10.

6.1 User Perspective

Article 26 of the DSA specifies key expectations for transparency from the perspective of end-users. The overarching goal of ad explanations is to enable individuals to understand and, where possible, control the use of their data. To this end, user-facing ad explanations should be: (1) **granular**; (2) **attributable**, i.e., clarifying whether targeting decisions stem from advertiser inputs or platform inferences; (3) **source-clear**, i.e., specifying the origin of the data used for targeting; and (4) **accurate**, faithfully reflecting the criteria that shaped ad delivery.

6.1.1 Granularity. Granularity refers to the level of detail provided in the ad explanation texts about the targeting reasons used to deliver a given ad. In the context of transparency regulation, including the DSA's requirement to provide *"meaningful information"* [1], granularity ensures that users are well-informed about the reasons for receiving an ad[5]. To evaluate granularity, we examine whether explanations specify: (1) the targeting categories applied (e.g., age, location, interests, custom audiences), and (2) the corresponding values (e.g., "aged 18–24," "based in Germany"). On this basis, we classify explanations into two levels:

Broad: Explanations provide only vague or generic information, such as "because of your interests," "time of day," or "based on your age," without specifying concrete values.
Precise: Explanations specify both categories and values (e.g., "aged 18–24 and located in Germany").

6.1.2 Attribution of Targeting and Delivery Choices. Attribution concerns whether explanations clearly identify who is responsible for the targeting criteria that led to ad delivery—whether advertisers, platforms, or a combination of both[1,5]. This helps users exercise control over how their data is used. We classify attribution in explanations into three categories:

[5]Recital (68) DSA "[..]Such explanations should include information on the method used for presenting the advertisement, for example whether it is contextual or other type of advertising, and, where applicable, the main profiling criteria used; it should also inform the recipient about any means available for them to change such criteria." [15]

Table 1: Explanation text classes with examples of regex/keywords used for categorization.

Targeting Class	Description	Regex/Keywords Examples
Behavioral	Targeting inferred from a user's activity (e.g., visits, views) or declared interests.	"visited (?:websites|pages) (?:[\w\s]+) topic(s?)"
Contextual	Targeting based on page/post content, video, search terms, topics, or keywords.	"near (?:posts|tweets) (?:with|containing) the keyword"
Custom Audience	Advertisers reaching users based on pre-existing customer data (e.g., hashed emails).	"hashed list used by (?:[\w\s]+)"
Demographic	Attributes such as age, gender, education, employment, language, or family status.	"(?:aged|age|ages) \d+[-+]?\d*"
Lookalike	Audiences built from similarity to existing audiences/customers.	"similar to (?:people|customers|existing customers)"
Retargeting	Targeting based on a user's past interactions with an advertiser (on-platform or external).	"(?:interacted|visited)(?: with| the)? (?:account|page|\website|products) (of|from) {*advertiser*}"
Time & Location	Mentions of time/location, or implicit references from device/IP address data.	"(?:current|main) (?:place|city|location|residence)"

Table 2: Distribution of ad explanation texts across targeting categories on Facebook, Instagram, YouTube, and X.

	Share of ad explanations			
Targeting Category	**Facebook**	**Instagram**	**YouTube**	**X**
Behavioral	7.86%	33.22%	12.66%	9.31%
Contextual	0.0%	0.0%	1.12%	0.28%
Custom Audience	0.49%	0.14%	0.65%	6.21%
Demographic	44.81%	33.32%	8.90%	30.48%
Lookalike	1.04%	0.55%	4.67%	8.31%
Retargeting	1.55%	0.61%	0.04%	5.85%
Time & Location	44.24%	32.16%	69.33%	38.42%
Other	0.0%	0.0%	2.63%	1.14%

Advertiser Choice: Delivery is explicitly tied to criteria selected by the advertiser (e.g., "The advertiser chose to target people aged 25–40 interested in skincare").
Platform Algorithm: Delivery is attributed to the platform's automated systems, inferences, or optimization processes (e.g., "You are seeing this ad because our systems identified you as likely interested in technology, based on your recent activity").
Ambiguous: Explanations use vague language that conflates advertiser and platform roles, obscuring accountability (e.g., "This ad is based on your interests").

6.1.3 Disclosure of Data Sources. Transparency also requires clarity about where the data originates. Explanations should identify whether targeting information derives from user-provided data, observed platform activity, off-platform tracking, or advertiser uploads. Such disclosure is critical to help individuals understand—and potentially contest—the pathways through which their data enters advertising systems[5]. We distinguish four categories of data sources:

User-Provided Data: Information explicitly supplied by users (e.g., profile age, gender, stated location).
On-Platform Activity: Signals inferred from user interactions within the platform (e.g., clicks, engagement, viewed content).
Off-Platform Activity: Data collected via external tracking mechanisms such as pixels, cookies, or third-party brokers.
Advertiser-Provided Data: Customer data uploaded by advertisers (e.g., hashed emails or phone numbers).

6.1.4 Accuracy. Explanations should faithfully reflect the actual parameters and logic used to determine why the user was shown a specific advertisement1. Targeting reasons that materially contribute to ad delivery should not be omitted; conversely, an attribute that did not influence delivery should not be presented

6.2 Auditor Perspective

Article 39 of the DSA requires very large online platforms (VLOPs) to provide public repositories of the advertisements running on their services. From an auditor's perspective—whether researcher, regulator, or journalist—such repositories should be: (1) **granular**, (2) **attributable**, and (3) **accurate**.

6.2.1 Granularity: Repositories should disclose the full set of targeting categories selected by advertisers (e.g., demographic attributes, customer lists, interests) and the associated values (e.g., specific age ranges, selected interests, or targeted regions). This should be done without revealing any personally identifiable information (PII) about user[3].

6.2.2 Attribution: Recital 95 of the DSA also specifies that repositories must include both targeting and delivery criteria [6]. This implies that repositories should clearly attribute responsibility for the criteria used in audience selection—distinguishing between advertiser choices and platform-level interventions such as algorithmic optimization.

6.2.3 Accuracy: Finally, repositories should provide a comprehensive and precise account of both the targeting criteria specified by advertisers and the delivery criteria applied by the platform when serving ads[7].

7 Evaluating User-facing Ad Explanations

In this section, we evaluate the granularity, the attribution of targeting and delivery choices, and the disclosure of data sources in ad explanations.

7.1 Methodology

After each explanation is classified by targeting category (as described in Section 4.1), we systematically benchmark the expressed

[6]DSA Recital 95 "Repositories should include the content of advertisements[..]. This information should include both information about targeting criteria and delivery criteria[..]" [15]
[7]DSA Article 39(1) (bolded by authors)"[..]They shall ensure that the repository does not contain any personal data of the recipients of the service to whom the advertisement was or could have been presented, and shall make reasonable efforts to ensure that the information is **accurate** and **complete**." [14]

transparency of user-facing explanations across platforms against DSA requirements, by determining the degree to which it satisfies the dimensions in Section 6.1.

We operationalize this evaluation using a **rule-based text classification** approach. For each transparency dimension, we construct regular expressions to identify linguistic markers in the explanation texts (e.g., references to data origins such as *"from your profile,"* or attribution cues such as *"the advertiser chose"*). Each WAIST explanation is automatically evaluated against these patterns, and the resulting labels are aggregated by targeting category and platform. The results are summarized in Table 3.

For Instagram, YouTube, and X, the WAIST data consists of user-facing strings or short phrases, which makes automated text evaluation straightforward. In contrast, Facebook's WAIST data was obtained by intercepting network requests and extracting structured JSON objects rather than rendered HTML. These objects contain internal attribute fields that the platform uses to generate WAIST notices dynamically. To evaluate Facebook explanations consistently with the other platforms, we mapped each attribute value in the JSON data to its corresponding user-facing phrase by reproducing the HTML rendering process. The resulting text strings were then processed using the same pattern-matching and classification method as applied to the other platforms. We provide the Jupyter Notebooks used for our analyses and the CSV files for all analyzed explanation texts [8].

7.2 Results

Granularity. Facebook, Instagram, and X have the highest shares of precise ad explanations. We find that all demographic, location, and behavioral categories are disclosed with their values in Facebook and Instagram, and almost all of them in X. In total, 98.58% of ad explanation texts were classified as precise on Facebook, 99.6% on Instagram, and 76.8% on X. The only exceptions are explanations for custom and lookalike audience targeting, which do not reveal the identifier type (such as email, name or phone number) or seed audiences (such as websites where the user was tracked) were used. By contrast, YouTube relies almost exclusively on broad explanations. 98.9% of all explanation texts cite only the main targeting form (e.g., "Your activity"), and do not detail the exact attributes used.

Attribution. Across platforms, attribution in ad explanations generally emphasizes advertiser choices. On Facebook and Instagram, the interface explicitly separates the advertiser-selected criteria from platform-inferred signals (see Figure 1). On Instagram, 56% of explanations appear under *"Advertiser choices"*, while 33% appear as platform inferences from user activity. On X, attribution is even more advertiser-centric. Explanation texts reference advertisers almost exclusively (e.g., "*The advertiser wants to reach people who are older than 21 years and are located in Germany*"). In our dataset, 88.6% reference advertisers, and none reference algorithmic interference in targeting. By contrast, YouTube provides little explicit attribution. Most explanations (84%) omit any reference to who determined the targeting criteria, using generic formulations such as "Websites that you've visited". Only 11% attribute targeting to

[8] https://osf.io/7nxj6/overview?view_only=b49c7a833965443bb8f4016c74b82a90

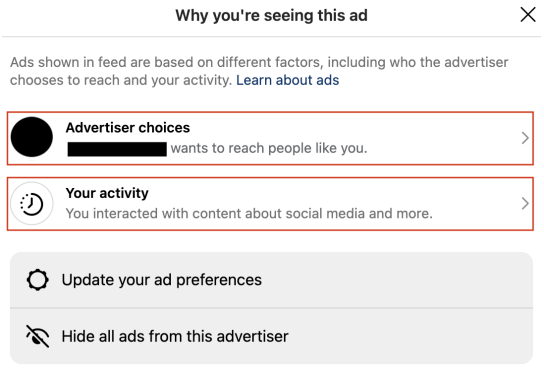

Figure 1: Example screenshot of the *"Why am I seeing this ad?"* user interface on Instagram. Details of attributes selected by advertisers are provided separately when the user interacts with "Advertiser choices" panel.

Google's algorithmic systems (e.g., "Google's estimation of your areas of interest"), particularly in behavioral ads, where such phrasing accounted for 95.6% of cases.

Disclosure. On Facebook, we analyze only explanations for which the data source could be clearly mapped to a targeting attribute (52.3% of all explanations). 84.3% of the matched ad explanations explicitly stated the data source used to deliver the ad. We find that 82.3% of them refer to user data (e.g., *"People who have indicated Single as their relationship status on Facebook"*). On Instagram, 72% of explanations explicitly mention the data used for targeting. Most of these (86.8%) refer to on-platform interactions, such as reactions to posts, interactions with content and accounts within Meta products (see Figure 1). Behavioral and remarketing ads are almost exclusively described as derived from such on-platform activity (e.g., *"You have interacted with posts, stories, reels, and pages on the following topics: Household goods"*, *"People who have interacted with the Instagram account of [advertiser]"*). In contrast, YouTube and X rarely provide data sources. Around 85% of explanations on both platforms lack any reference to where the targeting data comes from. On YouTube, explanations typically make vague references to *"your activity while signed in to Google"* (13.7%) or *"websites you've visited"* (0.4%), without specifying which activities or sites are involved. No distinction is made between Google's own tracking systems (e.g., remarketing tags) and data collected from external sources. On X, advertiser-provided data is mentioned in roughly 12% of explanations. At the same time, only 2.66% and 0.35% refer to off-platform and on-platform activity, respectively, with brief and vague statements such as "your recent web activity." Most explanations (85%)—especially those referencing behavioral targeting (71.4%)—omit data source information entirely, using generic wording like *"[advertiser] wants to reach people interested in data centers."*

7.3 Discussion

Our findings show that the current ecosystem of user-level transparency is uneven across platforms.

Table 3: Distribution of ad explanation texts by targeting category, analyzed by granularity, attribution of targeting and delivery choices, and data source disclosure across Facebook, Instagram, YouTube, and X.

Category	Platform	Granularity		Attribution			Data Source				
		Broad	Precise	Platform	Advertiser	Ambiguous	User Data	On-Platform	Off-Platform	Advertiser	Not Provided
Behavioral	Facebook	0.0%	100%	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Instagram	0.0%	100%	96.9%	3.1%	0.0%	0.0%	95.7%	1.2%	0.0%	3.1%
	YouTube	100%	0.0%	96.5%	0.0%	3.5%	0.0%	91.7%	3.5%	0.0%	4.7%
	X	28.6%	71.4%	0.0%	71.4%	28.6%	0.0%	0.0%	28.6%	0.0%	71.4%
Contextual	Facebook	0.0%	0.0%	0.0%	0.0%	0.0%	–	–	–	–	–
	Instagram	0.0%	0.0%	0.0%	0.0%	0.0%	–	–	–	–	–
	YouTube	0.0%	100%	0.0%	0.0%	100%	–	–	–	–	–
	X	0.0%	100%	0.0%	0.0%	100%	–	–	–	–	–
Custom Audiences	Facebook	100%	0.0%	0.0%	100%	0.0%	0.0%	0.0%	0.0%	100%	0.0%
	Instagram	100%	0.0%	0.0%	100%	0.0%	0.0%	0.0%	0.0%	100%	0.0%
	YouTube	35.5%	64.5%	0.0%	100.0%	0.0%	0.0%	0.0%	0.0%	100%	0.0%
	X	100%	0.0%	0.0%	100%	0.0%	0.0%	0.0%	0.0%	100%	0.0%
Demographic	Facebook	0.0%	100%	–	100%	0.0%	85.8%	0.0%	0.0%	0.0%	14.2%
	Instagram	0.0%	100%	–	100%	0.0%	16.2%	0.0%	0.0%	0.0%	83.8%
	YouTube	100%	0.0%	–	0.0%	100%	0.0%	0.0%	0.0%	0.0%	100%
	X	0.0%	100%	–	100%	0.0%	0.0%	0.0%	0.0%	0.0%	100%
Lookalike	Facebook	100%	0.0%	0.0%	100%	0.0%	0.0%	0.0%	0.0%	0.0%	100%
	Instagram	100%	0.0%	0.0%	100%	0.0%	0%	0%	0.0%	0.0%	100%
	YouTube	100%	0.0%	0.0%	0.0%	100%	0.0%	100%	0.0%	0.0%	0.0%
	X	100%	0.0%	0.0%	0.0%	100%	0.0%	0.0%	0.0%	0.0%	100%
Retargeting	Facebook	0.0%	100%	0.0%	100%	0.0%	0.0%	65.0%	35.0%	0.0%	0.0%
	Instagram	0.0%	100%	0.0%	100%	0.0%	0.0%	75.4%	24.6%	0.0%	0.0%
	YouTube	0.0%	100%	0.0%	0.0%	100%	0.0%	50%	50%	0.0%	0.0%
	X	98.8%	1.2%	0.0%	98.8%	1.2%	0.0%	1.2%	0.0%	98.8%	0.0%
Time & Location	Facebook	0.0%	100%	0.0%	100%	0.0%	N/A	N/A	N/A	N/A	N/A
	Instagram	0.0%	100%	0.0%	100%	0.0%	16.3%	50.0%	0.0%	0.0%	33.7%
	YouTube	100%	0.0%	1.0%	0.0%	99.0%	0.0%	0.0%	0.0%	0.0%	100%
	X	0.0%	100%	0.0%	100%	0.0%	0.0%	0.0%	0.0%	0.0%	100%

First, the variation in granularity suggests that the *"meaningful information"* standard is being interpreted inconsistently, leaving users with very different levels of insight into why they are shown particular ads across platforms. Explanations on Facebook and Instagram are relatively detailed, while on YouTube, they remain almost entirely broad and generic. X falls in between, providing partial disclosure of behavioral data but offering only ambiguous treatment of specific targeting forms like retargeting, custom, and lookalike audiences.

Additionally, we find that users are rarely given a clear account of whether the criteria used to show them the ads were selected by advertisers or due to platform-driven inferences and optimizations. Meta's platforms (Facebook and Instagram) provide a clear indication of targeting decisions through the WAIST interface and explanation text (e.g., "*advertiser identified a group of customers that it wanted to use to find similar customers. Meta technologies used that information to find potential customers.*"), whereas X emphasizes advertiser choice while downplaying its own algorithmic contribution. YouTube offers little attribution, with most explanations failing to disclose any party's involvement in decision-making. This obscures accountability in the broader advertising ecosystem.

Finally, the disclosure of data sources is limited. Across all platforms, we find little to no direct reference to off-platform activity, despite the evidence that these platforms operate as large-scale cross-site trackers [36]. Even when categories such as custom or lookalike audiences are disclosed, users are not told how those audiences were constructed, what identifiers were used, or whether off-platform tracking contributed to their inclusion. This lack of specificity constrains users' ability to assess the scope of data use or exercise control over targeting.

Together, these shortcomings point to a transparency landscape that is still heavily **platform-curated**. We find that Instagram has overall the best implementation of user-level transparency in terms of level of details (granularity of ad explanations), the disclosure of targeting and delivery decisions, and transparency of data sources.

Compliance. The observed practices raise questions about compliance with the Digital Services Act's transparency obligations in Article 26. The broad and incomplete nature of many targeting explanations—especially those that fail to specify the concrete data sources or targeting parameters involved—suggests that platforms may not be meeting the standard of meaningfulness envisioned by the regulation. Merely indicating that an ad is shown *"based on your activity"* or *"because you might be interested in X"* does not enable a reasonable user to understand or contest the underlying data processing, undermining both the right to information and the DSA's aim of user empowerment[5]. Similarly, limited disclosure of whether targeting relies on first-party, inferred, or third-party data conflicts with the DSA's expectation of clarity regarding data provenance[1]. In this sense, current implementations appear to deliver only a formal form of compliance, without providing the interpretability and accountability that the DSA legally requires.

8 Accuracy of User-facing Explanations on YouTube

Among the platforms we study, YouTube stands out for the unusual character of its ad explanations. While Facebook, Instagram, and X provide at least reasonably consistent disclosures—even if they fall short on some criteria—YouTube's explanations often appear to be vague, and repetitive (e.g., *"Google's estimation of your interests"*, *"Google's estimation of areas of your interest"*, *"Google's estimation of your interests based on your activity"* which could all be shown at once), which raises questions about whether they are dynamically generated.

Furthermore, prior work by Medjkoune et al. [29] documents inconsistencies in YouTube's ad explanations, such as failing to disclose that advertisers had explicitly selected placements (specific videos), to show their ads. These findings raise broader doubts about the correctness of explanations and whether they faithfully represent the targeting criteria actually used. In this section, we conduct a targeted analysis of the accuracy of YouTube's explanations.

Methodology. To conduct our analysis, we implement an experiment design where user profiles with various personalization settings are **served the same targeted ad**. We would then compare the explanations shown to each profile to find possible inconsistencies. The targeting strategy used must reliably deliver the same ads to the profiles, with minimal to no variation. It cannot rely on mechanisms such as interest-based targeting, where the absence of a disclosure could not be clearly attributed to YouTube's lack of transparency. For this reason, we selected **search terms**. Search-term targeting is a common feature on YouTube and can be consistently triggered across multiple profiles under controlled conditions.

We created four distinct profiles on Google and YouTube: **(1)** A non-logged-in user; **(2)** A logged-in user with ad personalization disabled; **(3)** a logged-in user with personalization enabled and YouTube search history disabled; and **(4)** a user with personalization and search history enabled. All profiles were newly created and lacked prior watch history, ad interaction, or broader Google account activity.

Each profile executed a series of ad-relevant searches (e.g., "Travel to [location]", "Online Banking", "Rent a car"), after which we crawled all ads shown, capturing advertiser details, campaign IDs, ad URLs, and ad explanations.

Across 60 search queries, we identified 11 unique ads that were shown across all profiles (same advertiser, same campaign, and identical URL parameters). Ad explanations shown to the profile with enabled personalization and search history consistently cited *"Search Terms"* as a targeting rationale. However, explanations shown to the no profile, and the disabled personalization and history users, all omitted it, only citing *"time of day and general location"*.

Discussion. Overall, it is difficult to verify the accuracy of ad explanations on any platform from a third-party standpoint, as users and auditors must rely on the information that platforms choose to disclose. In this experiment, we focused on observable discrepancies in the targeting explanations provided to profiles with varying data-processing settings. Our results suggest that the platform relied on specific targeting signals, such as search terms or keywords, but selectively disclosed these details depending on the profile type. This practice erodes trust in explanations as a tool for targeting transparency and undermines the DSA's requirement to provide *"meaningful information."*

9 Longitudinal Analysis of Facebook's Ad Explanations

We analyze **38,011 Facebook WAIST objects** collected between September 2017 and March 2025 to examine how the platform's ad explanations evolved. The analytical procedure follows the same approach described in Section 7. The data were matched to explanation texts corresponding to each attribute based on prior work and Facebook product announcements [2, 3, 7, 26].

Our longitudinal analysis serves two goals: (1) to trace changes in the targeting criteria disclosed to users, and (2) to benchmark transparency against our evaluative dimensions—granularity, attribution, and data-source disclosure—across distinct regulatory and product phases.

The dataset reveals a clear evolution in how targeting information is represented to users, aligning with major regulatory milestones and platform redesigns. We identify three distinct phases: (1) **the pre-GDPR period (2017–2019)**, (2) **the post-GDPR and privacy-policy expansion phase (2019–2022)**, and (3) **the post-DSA redesign (2023–2025)**. The following describes how user-facing transparency developed across these phases according to our criteria. Table 7 summarizes our assessment of each targeting attribute and category over time.

Our observations reflect what was available to our German user panel, and may not align precisely with Facebook's global rollout or official announcements. Some categories appear in our dataset later than in Meta's public announcements, likely due to regional differences, staggered feature deployment, or data collection gaps. Accordingly, the evolution reported here should be interpreted as empirical evidence drawn from user data, not an exhaustive chronology of Facebook's internal system changes.

2017–2019 (Pre-GDPR/CCPA). During this period, ad explanations mirrored the structure of Facebook's Ads Manager and disclosed straightforward targeting parameters—age, gender, location, language, and interests—often with explicit values (e.g., "*people aged 25–40 living in Berlin*"). All explanations typically attributed targeting to advertiser choices (e.g., "*One reason you're seeing this ad is that [advertiser] wants to reach*") and referenced user-provided data in 60% of cases or on-platform data in 30%, such as profile fields or page likes. Custom and lookalike audiences also appeared but were described generically, without any disclosure of the underlying data used to form those segments.

2019–2022 (Post-GDPR). Following the GDPR's enforcement and Meta's update in 2019 to WAIST notices [7], explanations incorporated behavioral and engagement-based criteria such as page interactions, app usage, and video views, alongside demographic fields like employer or education. This increased descriptive richness but also introduced some ambiguity. For instance, explanations referring to user interactions such as "*based on activity such as liking pages or clicking on ads*" blurred the distinction between advertiser targeting and Meta's own optimization systems. Finally,

Table 4: Longitudinal analysis results on Facebook from September 2017 to March 2025.

Category	Time period	Granularity		Attribution		Data Source					
		Broad	Precise	Advertiser	N/A	User Data	On-Platform	Off-Platform	Advertiser	Not Provided	N/A
Behavioral	2017-2025	0.0%	100%	98.3%	1.7%	0.0%	88.0%	0.0%	0.0%	2.3%	9.6%
Custom Audience	2017-2025	100%	0.0%	100%	-	0.0%	0.0%	0.0%	100%	0.0%	-
Demographic	2017-2025	0.0%	100%	100%	-	96.5%	0.0%	0.0%	0.0%	3.5%	-
Lookalike	2017-2025	100%	0.0%	100%	-	0.0%	0.0%	0.0%	0.0%	100%	-
Retargeting	2017-2025	0.0%	100%	100%	-	0.0%	21.5%	7.5%	71.0%	0.0%	-
Time & Location	2017-2025	0.0%	100%	100%	-	74.7%	0.0%	0.0%	0.0%	0.0%	25.3%

data-source statements became less specific, most of the explanations referred to user-stated data (83.7%) and often omitted whether information came from on-platform signals, advertiser uploads, or off-site tracking.

2023–2025 (Post-DSA). Coinciding with the rollout of the DSA and Meta's Transparency Center, explanations underwent a major structural change. The JSON schema is split into two components: (1) *advertiser-defined parameters* (demographics, interests, audience lists) and (2) *aggregated topic-based explanations summarizing inferred interests on and off Meta technologies* (e.g., "*You have interacted with ads about travel.*") (see Figure 1 and 2). New formats such as Collaborative Ads, lead-generation campaigns, and offline conversions via the Conversion API began to appear. This redesign clarified attribution by separating advertiser inputs from algorithmic delivery, but reduced granularity and data-source visibility, as explicit values (e.g., employer, education, connections) were removed and data provenance was no longer specified.

Discussion. Over time, Facebook's ad explanations evolved from explicit, attribute-level disclosures to abstract, topic-based rationales. While Meta has improved **attribution clarity** by explicitly distinguishing advertiser-defined criteria from platform-level delivery optimization, both **granularity** and **data-source transparency** have declined as targeting has shifted toward algorithmic inference and aggregated reporting. As ad delivery becomes less dependent on discrete advertiser-selected attributes, user-facing explanations are increasingly difficult to map onto interpretable targeting categories. The resulting transparency falls short of meaningful interpretability, limiting the ability of users and auditors to understand or verify the data and logic underlying ad targeting decisions.

10 Evaluating Public Ad Repositories

In this section, we assess the implementations of ad repositories on Meta, Google, and X against the criteria defined in Section 6.2.

10.1 Methodology

Granularity. For each platform, we extracted and reviewed all targeting-related fields exposed through the public interface or API (e.g., audience types, demographics, locations, interests). We assessed the granularity of these fields by examining (1) whether they matched the categories and level of detail shown in user-facing explanations, and (2) whether they corresponded to the full set of targeting options available to advertisers within the platform's ad managers [25, 31, 43]. As with user-facing explanations, we evaluated the specificity of each field, whether they included only high-level categories (e.g., "demographics," "geography") or also the specific values applied by advertisers (e.g., "aged 18–24," "Berlin," "sports interest"), and their coverage.

Attribution. We assessed attribution by examining how each repository represents responsibility for targeting decisions. Specifically, we evaluated whether targeting parameters were labeled as advertiser-selected or whether the repository acknowledged the role of platform-level optimization systems in determining delivery.

Accuracy. We measure accuracy by analyzing the consistency between the targeting information disclosed to users and the targeting parameters reported in the ad repository for the same ad. Each user-facing ad explanation is linked to its corresponding repository record when a unique match can be determined using shared identifiers such as advert ID, advertiser name, and campaign metadata. Cases where no matching record exists or where the repository entry is incomplete are excluded from further analysis. Then, for each matched ad, we identify two types of discrepancies:

Missing category: A targeting category present in the user-facing explanation but absent from the repository entry.
Discrepant attribute: A mismatch between the values in the repository and those in the user-facing explanation. For example, when the repository specifies targeting users in Berlin, but the user-facing explanation lists Paris.

We only consider omissions relative to what is shown to the user as missing; the reverse (repository attributes absent from the user's explanation) is not treated as an inaccuracy, since different users may legitimately see different subsets of targeting attributes depending on delivery context. For each ad, we compute the proportion of attributes that fall into these discrepancy categories and summarize the results by targeting class and platform (see Table 5).

10.2 Meta's Ad Repository

We analyze *Meta's Ad Library*, which hosts ads (active and inactive) running across Meta platforms (Instagram, Facebook, Threads, Messenger, and their Audience Network[9]).

10.2.1 Granularity. Meta discloses only three targeting categories along with their values: age, gender and location. The repository omits all other targeting criteria. As a result, discrepancies across targeting types visible in user-facing explanations vs. in the repository cannot be comprehensively evaluated.

[9] Mobile apps that partner with Meta for running advertisements.

Table 5: Discrepancy analysis between targeting explanations presented to users and those found in the ad repositories.

Platform	Targeting Category	User-Facing	Repository	Missing Category %	Discrepant Attribute %
Google	Behavioral	79	108	**41.77%**	N/A
	Contextual	0	189	0%	N/A
	Custom Audience	1	131	0%	N/A
	Location	189	189	0%	N/A
	Lookalike	41	0	**100%**	N/A
	Time	131	0	**100%**	N/A
Twitter	Behavioral	257	382	**68.09%**	**95.56%**
	Contextual	9	991	0%	0%
	Custom Audience	440	77	**92.95%**	N/A
	Demographic	1064	1220	0.38%	0.28%
	Location	1427	1428	0%	0.21%
	Lookalike Audiences	432	14	**99.31%**	**92.68%**
	Retargeting	84	246	**67.86%**	N/A
	No Explanation Provided	15	0	N/A	N/A
Meta	Behavioral	422	0	100%	N/A
	Custom Audience	5	0	100%	N/A
	Demographic	580	580	0%	0%
	Location	565	565	0%	12.20%
	Lookalike	15	0	100%	N/A
	Retargeting	11	0	100%	N/A

10.2.2 Attribution. In terms of attribution, all selections are designated as advertiser choices, which is consistent with how ad campaigns are set up in the Ads Manager [31] (required attributes).

10.2.3 Accuracy. We link the ads **with valid advertisement IDs** in our dataset (September 2024 to March 2025) to their images in the repository. This yielded 181 out of 351 matches on Facebook (48% missing from the repository) and 399 out of 438 on Instagram (9% missing).

When comparing targeting details, we find that all age and gender values found in user-facing ad explanations match the repository. There were no missing or discrepant data. For location targeting, the repository was often more detailed than explanations (e.g., *"Bielefeld, Germany"* in the repository vs. *"The current place of residence that you have specified in your Facebook profile. People with their main location in Germany"* in user-facing explanation). This, however, appears to be the case only when targeted locations chosen by the advertiser are granular (such as a city or zip code rather than a country). We identify this as a discrepancy but do not consider it to be problematic. This was the case for 17% of matched ads on Instagram, and did not include Facebook, as our matching procedure of JSON attributes does not allow capturing such nuances.

10.3 Google's Ad Repository

We analyze Google's Transparency Center, which consolidates information about ads running across Google services (e.g., YouTube, Google Maps, Google Play).

10.3.1 Granularity. The repository discloses the use of five *high-level* targeting categories: *"Demographic info"*, *"Geographical locations"*, *"Contextual signals"*, *"Customer lists"*, and *"Topics of interest"*, without the specific values selected (e.g., exact age ranges, regions, or interests). These categories do not align with user-facing ad explanations (see Table 1), nor do they align with the options available to advertisers, as observed on the ad manager [23].

10.3.2 Attribution. All targeting criteria in the repository interface are presented as *advertiser-selected*. This raises the question of whether such criteria are systematically included as advertiser selections, regardless of whether algorithmic delivery played a role in their determination.

Experiment. To examine attribution correctness, we conducted a controlled experiment using *Google Ads Manager*. Acting as an advertiser, we created a YouTube campaign targeting users in France. We deliberately **selected no targeting parameters** other than location, which is a mandatory field, leaving all other settings (demographics, interests, keywords, devices, and languages) at their default, unrestricted state. We do, however, enable algorithmic optimization settings, which let the platform determine the relevant audience segments to whom the ads are shown.

Results: When the campaign appeared in the repository, it listed three targeting parameters: *"Geographical location"*, *"Demographical information"*, and *"Contextual signals"*, all attributed to the advertiser. The latter two **were not selected** during campaign configuration.

The inclusion of demographic and contextual categories may reflect Google's automated delivery optimization, which dynamically adjusts targeting to improve performance. However, the repository explicitly attributes these parameters to the advertiser, implying direct selection. If such parameters are determined algorithmically, their presentation as advertiser choices is misleading and raises concerns about attribution accuracy. If these targeting parameters are not in fact used by the platform, they constitute a clear discrepancy. Overall, this experiment suggests that Google's ad repository may misrepresent targeting parameters by, at best, conflating advertiser-specified criteria with platform-inferred signals, and, at worst, being fundamentally unreliable.

10.3.3 Accuracy. Because no direct identifiers link the ads in user-facing explanations to their corresponding images in Google's Transparency Center, we relied on the `video_id` embedded in

some ad URLs to establish connections. To do so, we scraped advertiser pages and searched for thumbnail URLs containing the same `video_id` values present in our dataset. This approach yielded 806 potential matches (23% of ads with a `video_id`). After excluding cases with multiple or ambiguous matches, only 189 ads (6% of our YouTube dataset) could be uniquely linked to a repository entry.

Because Google's repository discloses only broad targeting categories and no exact values, our accuracy assessment focuses solely on missing categories. We find that 41.7% of explanations citing behavioral (interest-based) targeting do not have a corresponding parameter in their repository entry; we treat Google's "Topics of interest" field as the closest match for this category. In addition, all lookalike and schedule (e.g., "Time of day") targeting is omitted, as no repository field captures this targeting strategy (100% of matched ads with this explanation).

10.4 X's Ad Repository

In this section, we analyze X's public ad repository, which provides information about ads running on the platform.

10.4.1 Granularity. Targeting information is disclosed in two fields: *"Targeted Segments"* and *"Excluded Targeting Segments"*. Both are unstructured token (attribute) lists that may represent demographics, locations, interests, or keywords, but without categorical labels. For example, *"computer games"* could denote an interest, a topic, or a keyword. The lack of documentation or consistent formatting complicates interpretation. This representation does not provide clear details about which targeting method (e.g., contextual or behavioral) is used.

10.4.2 Attribution. All targeting is attributed to *advertisers*. No information is provided about platform-level delivery or optimization. We attempted to conduct a similar controlled experiment on X (see Section 10.3) to verify the accuracy of attribution. However, this was not possible because obtaining advertiser verification on the platform took several months and was ultimately not approved within our study period.

10.4.3 Accuracy. To connect user-facing ads with entries in X's ad repository, we relied on the advertiser name provided in the ad explanations. For each ad, we retrieved the corresponding CSV file for that advertiser and campaign period, then extracted the Tweet ID from the promoted Tweet or media URL. Using these identifiers, we matched ads in our dataset to repository entries. In total, we identified advertisers with repository entries for 2,190 ads. Of these, 1,458 ads were successfully matched using Tweet IDs.

To conduct our analysis, we categorized the unstructured targeting tokens using a set of heuristics. Hashtags were treated as indicators of contextual targeting (i.e., keyword-based), while account names were interpreted as lookalike targeting (specifically, follower lookalike). Demographic, location-based, retargeting, and custom-audience targeting were generally identifiable through explicit references in the tokens (e.g., *"You might be seeing this ad because you are part of or similar to an audience from CodeRabbit."*, *"One possible reason you might be seeing this ad is that you resemble the followers of Steam."*, *"...wants to reach people older than 25 years and located here: Germany."*). To identify behavioral targeting (i.e., interest-based), we cross-referenced the attributes extracted from repository entries against X's Ads API taxonomy [45]. Tokens corresponding to entries in the API's interests category were classified as behavioral targeting.

We first assess missing targeting categories. We find substantial omissions: 68% of ads shown to users with behavioral (interest-based) explanations, 93% of ads referencing custom audiences, and virtually all explanations citing lookalike targeting (99%), did not mention these categories in the repository. Retargeting information was absent in 68% of cases. Only contextual, demographic, and location-based targeting categories matched between explanations and repository entries, with only a couple of omissions.

We also observe major qualitative discrepancies in two key categories—behavioral and lookalike audiences. These appear in two forms. First, the repository often contains generalized or loosely related tokens rather than the attributes shown to users. For instance, a user might see *"because you are interested in luxury"*, while the repository lists *luxury cruises*, *Louis Vuitton*, and *Gucci*. In other cases, the discrepancy is more severe: a user may be told *"because you are interested in data centers"*, while the repository instead records selected keywords such as *laptop computers* and *desktop*. These inconsistencies raise questions about whether algorithmic delivery optimization modifies targeting criteria and, if so, why such adjustments are not disclosed.

10.5 Discussion

Our analysis of ad repositories (Meta, Google, and X) reveals that none provides a complete or reliable account of the targeting parameters used in advertising delivery. While each platform makes claims of transparency, their repositories diverge substantially in coverage, structure, and compliance with the DSA's transparency requirements.

Granularity. None of the repositories examined provides a comprehensive representation of targeting criteria. Meta's Ad Library does not disclose any targeting parameters (except for age, gender and locations). Google's Transparency Center lists five broad categories—demographic information, geography, contextual signals, customer lists, and topics of interest—but these fields neither align with the full set of targeting options available to advertisers nor include the specific values (e.g., "aged 18–24," "located in Paris") that would make them interpretable. On X, the repository includes detailed targeting values, but the lack of structure—such as unlabeled tokens—makes it difficult to identify which targeting strategy each value represents. The absence of standardized fields across platforms severely limits the interpretability and comparability of targeting disclosures.

Attribution. The DSA explicitly requires ad repositories to disclose both advertiser selections and platform optimization choices[6], thereby allowing accountability for targeting decisions. In practice, no repository in our study makes this distinction. All three attribute targeting parameter selections to the advertiser, even in cases where our controlled experiments demonstrate that platform algorithms introduced additional criteria (see Section 10.3). This conflation of responsibility obscures the role of automated delivery systems and prevents auditors or regulators from determining who—advertiser or platform—is responsible when targeting raises legal, ethical, or

discriminatory concerns. In some cases, such as Google Ads, even advertisers are unaware of the full subset of attributes used by optimization systems to refine delivery.

Accuracy. Accuracy is the most crucial criterion for transparent repositories. Yet our analysis indicates substantial discrepancies between user-facing explanations and repository records. On Google, several targeting categories visible to users—such as behavioral signals, lookalike audiences, and scheduling (time)—are missing from public records. On X, as well, attributes reported in user-facing explanations do not always appear. These inconsistencies could stem from omissions in repositories, inaccuracies in explanations, or both, but either scenario undermines the integrity of the transparency mechanisms.

We also note that all three repositories present notable usability challenges. None provides stable identifiers that enable direct linking between user-facing ads and repository entries. On Google, we had to scrape advertiser pages to locate corresponding ads, as neither unique identifiers nor cross-references exist. On X, ad data can only be accessed in bulk via advertiser names, constraining systematic analysis. These usability barriers substantially reduce the feasibility of external auditing—the very purpose of the repositories under the DSA.

Compliance. The empirical evidence indicates widespread noncompliance with the DSA's transparency obligations outlined in Article 39. None of the repositories provides auditors with complete and accurate data access to understand how targeting decisions are made, nor do they accurately attribute responsibility between advertisers and platforms. Meta's decision to restrict targeting data to a separate political advertising dataset is incompatible with the DSA's requirement that *all* ads be included in the public repository. Google and X, while offering broader data access, impose significant procedural and technical barriers and exhibit inconsistencies that call the reliability of their disclosures into question. Even if the source of these inaccuracies cannot be definitively attributed—whether to the user-facing explanations or to the repositories themselves—both outcomes imply a breach of the DSA's transparency principles (Articles 26 and 39). Inaccurate or incomplete disclosures fail to meet the legal standard of reliability and completeness set out in Article 39 and erode trust in the mechanisms designed to enable external oversight.

11 Limitations

Data were collected through the Who Targets Me browser extension, which relies on scraping platform interfaces. As a result, some dynamically loaded elements may not have been captured, and frequent changes to platform code can affect the consistency of the data collection over time. This meant we had a few gaps in the dataset, and some metadata was missing for some entries. Additionally, the data we acquired for Facebook was in the form of structured JSON objects that listed targeting attribute values, rather than textual explanations shown to users. To analyze this data, we matched the attributes with their corresponding textual elements across time. However, not all data were successfully matched, as some attributes appearing in our dataset did not seem to have a corresponding textual explanation. Finally, when assessing repository accuracy, we matched observed ads to their corresponding repository entries. For some ads, no matching entry could be found. Because the study was conducted within one year of the ads' delivery, this cannot be explained by the regulatory one-year retention period. While we could not determine the cause, this may indicate that repositories are occasionally incomplete or that ad entries are removed before the one-year deadline. As a result, we were unable to analyze all ads with identifiers for Meta and X.

12 Conclusion

This paper presented a multi-platform analysis of advertising transparency in the post-DSA environment, focusing on user-facing explanations and public ad repositories across Facebook, Instagram, YouTube, and X. Our results show that transparency remains fragmented and incomplete. Explanations often simplify or omit targeting details, limiting user understanding and control. While ad repository analysis shows discrepancies between what users are told and what repositories disclose, which highlight major obstacles to meaningful auditing. These findings show that no platform currently delivers the level of transparency envisioned by the DSA, either for users or for external auditors.

Acknowledgments

This research was supported by EU Grant No. 101041223. We thank the anonymous reviewers for their helpful and constructive feedback. We are also grateful to the *Who Targets Me* team for providing access to the data used in this study.

The authors used AI-based tools, including OpenAI ChatGPT, Gemini, and Grammarly, to assist with correcting typographical errors, grammatical issues, and awkward phrasing throughout the manuscript.

References

[1] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through optimization: How Facebook's Ad delivery can lead to biased outcomes. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–30.

[2] Athanasios Andreou, Marcio Silva, Fabrício Benevenuto, Oana Goga, Patrick Loiseau, and Alan Mislove. 2019. Measuring the Facebook Advertising Ecosystem. In *NDSS 2019 - Proceedings of the Network and Distributed System Security Symposium*. San Diego, United States, 1–15. https://doi.org/10.14722/ndss.2019.23280

[3] Athanasios Andreou, Giridhari Venkatadri, Oana Goga, Krishna P Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Investigating Ad Transparency Mechanisms in Social Media: A Case Study of Facebook's Explanations. In *NDSS 2018 - Network and Distributed System Security Symposium*. San Diego, United States, 1–15. https://doi.org/10.14722/ndss.2018.23204

[4] Muhammad Ahmad Bashir, Umar Farooq, Maryam Shahid, Muhammad Fareed Zaffar, and Christo Wilson. 2019. Quantity vs. Quality: Evaluating User Interest Profiles Using Ad Preference Managers. In *Proceedings 2019 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. https://doi.org/10.14722/ndss.2019.23392

[5] Muhammad Ahmad Bashir and Christo Wilson. 2018. Diffusion of User Tracking Data in the Online Advertising Ecosystem. *Proceedings on Privacy Enhancing Technologies* 2018, 4 (Oct. 2018), 85–103. https://doi.org/10.1515/popets-2018-0033

[6] bdarwell. 2014. *Making Ads Better and Giving People More Control Over the Ads They See.* https://about.fb.com/news/2014/06/making-ads-better-and-giving-people-more-control-over-the-ads-they-see/

[7] bdarwell. 2019. *Why Am I Seeing This? We Have an Answer for You.* https://about.fb.com/news/2019/03/why-am-i-seeing-this/

[8] X Business. 2025. Optimized Targeting. https://business.x.com/en/help/campaign-setup/campaign-targeting/optimized-targeting#3

[9] Jane Castleman and Aleksandra Korolova. 2025. Why am I Still Seeing This: Measuring the Effectiveness of Ad Controls and Explanations in AI-Mediated Ad Targeting Systems. In *Proceedings of the 2024 AAAI/ACM Conference on AI, Ethics, and Society*. AAAI Press, 255–266. https://dl.acm.org/doi/pdf/10.5555/3716662.3716685
[10] Salim Chouaki, Islem Bouzenia, Oana Goga, and Beatrice Roussillon. 2022. Exploring the Online Micro-targeting Practices of Small, Medium, and Large Businesses. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 1–23. https://doi.org/10.1145/3555103 Publisher: Association for Computing Machinery (ACM).
[11] John Cook, Rishab Nithyanand, and Zubair Shafiq. 2020. Inferring Tracker-Advertiser Relationships in the Online Advertising Ecosystem using Header Bidding. *Proceedings on Privacy Enhancing Technologies* 2020, 1 (Jan. 2020), 65–82. https://doi.org/10.2478/popets-2020-0005
[12] Laura Edelson, Tobias Lauinger, and Damon McCoy. 2020. A Security Analysis of the Facebook Ad Library. In *2020 IEEE Symposium on Security and Privacy (SP)*. San Francisco, CA, USA, 661–678. https://doi.org/10.1109/SP40000.2020.00084 ISSN: 2375-1207.
[13] European Parliament and Council of the European Union. 2022. *Article 26, the Digital Services Act (DSA)*. https://www.eu-digital-services-act.com/Digital_Services_Act_Article_26.html
[14] European Parliament and Council of the European Union. 2022. *Article 39, the Digital Services Act (DSA)*. https://www.eu-digital-services-act.com/Digital_Services_Act_Article_39.html
[15] European Parliament and Council of the European Union. 2022. *Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act)*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022R2065 OJ L 277, 27.10.2022.
[16] Facebook. 2025. Ad Preferences | Facebook. https://www.facebook.com/adpreferences/ad_settings/
[17] Mozilla Foundation. 2024. Full Disclosure: Stress testing tech platforms' ad repositories. https://assets.mofoprod.net/network/documents/Full_Disclosure_Stress_Testing_Tech_Platforms_Ad_Repositories_3FepU2u.pdf
[18] Eleni Gkiouzepi, Athanasios Andreou, Oana Goga, and Patrick Loiseau. 2023. Collaborative Ad Transparency: Promises and Limitations. In *2023 IEEE Symposium on Security and Privacy (SP)*. 2639–2657. https://doi.org/10.1109/SP46215.2023.10179448 ISSN: 2375-1207.
[19] Google. 2025. About audience expansion - Google Ads Help. https://support.google.com/google-ads/answer/9496986?hl=en
[20] Google. 2025. About optimization - Google Ad Manager Help. https://support.google.com/admanager/answer/184094?hl=en
[21] Google. 2025. *Ad Controls and Personalization Settings - Google Safety Center*. https://safety.google/privacy/ads-and-data/
[22] Google. 2025. Contextual targeting - Google Ads Help. https://support.google.com/google-ads/answer/1726458?hl=en
[23] Google. 2025. *Google Ads*. https://business.google.com/en-all/google-ads/
[24] Google. 2025. My Ad Center. https://myadcenter.google.com/home?sasb=true&ref=ad-settings
[25] Google. 2025. Reach Customers & Sell More with Online Advertising - Google Ads. https://business.google.com/en-all/google-ads/
[26] jdamiani. 2023. *Increasing Our Ads Transparency*. https://about.fb.com/news/2023/02/increasing-our-ads-transparency/
[27] Levi Kaplan, Nicole Gerzon, Alan Mislove, and Piotr Sapiezynski. 2022. Measurement and analysis of implied identity in ad delivery optimization. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 195–209.
[28] Who Targets Me. 2025. *About Who Targets Me – Who Targets Me*. https://whotargets.me/en/about/
[29] Tinhinane Medjkoune, Oana Goga, and Juliette Senechal. 2023. Marketing to children through online targeted advertising: Targeting mechanisms and legal aspects. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 180–194.
[30] Meta. 2025. *About Advantage+ audience*. https://en-gb.facebook.com/business/help/273363992030035
[31] Meta. 2025. Ads Manager - Guided Onboarding. https://adsmanager.facebook.com/adsmanager
[32] Meta. 2025. How Does Facebook Use Machine Learning to Deliver Ads? | Meta for Business. https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads
[33] Meta. 2025. *Why am I seeing ads from an advertiser on Facebook? | Facebook Help Centre*. https://www.facebook.com/help/794535777607370/
[34] Meta. 2025. *Why am I seeing ads from an advertiser on Instagram? | Instagram Help Center*. https://help.instagram.com/609473930427331
[35] Mozilla. 2019. Facebook and Google: This is What an Effective Ad Archive API Looks Like. https://blog.mozilla.org/en/mozilla/facebook-and-google-this-is-what-an-effective-ad-archive-api-looks-like/
[36] Sebastian Schelter and Jérôme Kunegis. 2016. Tracking the Trackers: A Large-Scale Analysis of Embedded Web Trackers. *Proceedings of the International AAAI Conference on Web and Social Media* 10, 1 (2016), 679–682. https://doi.org/10.1609/icwsm.v10i1.14769
[37] stuartlerman. 2022. *Introducing New Automation Tools to Increase Sales and Drive Growth*. https://about.fb.com/news/2022/08/introducing-new-automation-tools-to-increase-sales-and-drive-growth/
[38] Michael Carl Tschantz, Serge Egelman, Jaeyoung Choi, Nicholas Weaver, and Gerald Friedland. 2018. The accuracy of the demographic inferences shown on Google's Ad Settings. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*. 33–41.
[39] Giridhari Venkatadri, Athanasios Andreou, Yabing Liu, Alan Mislove, Krishna P. Gummadi, Patrick Loiseau, and Oana Goga. 2018. Privacy Risks with Facebook's PII-Based Targeting: Auditing a Data Broker's Advertising Interface. In *2018 IEEE Symposium on Security and Privacy (SP)*. 89–107. https://doi.org/10.1109/SP.2018.00014 ISSN: 2375-1207.
[40] Miranda Wei, Madison Stamos, Sophie Veys, Nathan Reitinger, Justin Goodman, Margot Herman, Dorota Filipczuk, Ben Weinshel, Michelle L. Mazurek, and Blase Ur. 2020. What Twitter Knows: Characterizing Ad Targeting Practices, User Perceptions, and Ad Explanations Through Users' Own Twitter Data. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 145–162. https://www.usenix.org/system/files/sec20-wei.pdf
[41] Ben Weinshel, Miranda Wei, Mainack Mondal, Euirim Choi, Shawn Shan, Claire Dolin, Michelle L. Mazurek, and Blase Ur. 2019. Oh, the Places You've Been! User Reactions to Longitudinal Transparency About Third-Party Web Tracking and Inferencing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, London United Kingdom, 149–166. https://doi.org/10.1145/3319535.3363200
[42] X. 2022. Your privacy options for personalized ads | X Help. https://help.x.com/en/safety-and-security/privacy-controls-for-tailored-ads
[43] X. 2025. Advertising. https://business.x.com/en/advertising
[44] X. 2025. X Ads targeting. https://business.x.com/en/advertising/targeting
[45] X Developer Platform. 2025. *Targeting Options*. https://developer.x.com/en/docs/x-ads-api/campaign-management/api-reference/targeting-options

A Ad Explanation Examples

Table 6 shows examples of ad explanation texts found in platform WAIST notices on Instagram, YouTube and X (ex. Twitter).

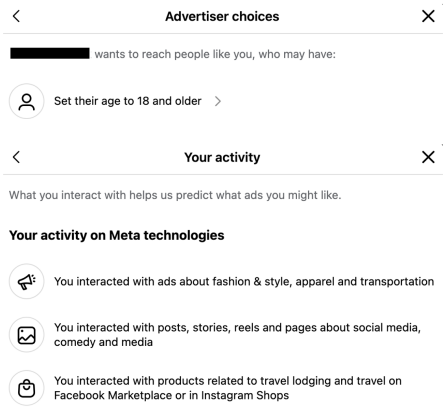

Figure 2: Screenshot of WAIST notice on Instagram: Advertiser choices and Your activity (platform algorithm).

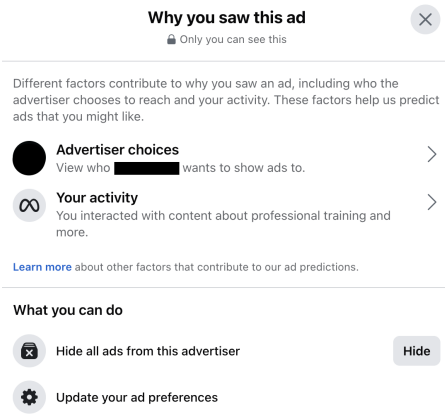

Figure 3: Screenshot of WAIST notice on Facebook.

A.1 Facebook WAIST

Below is a sample of Facebook WAIST data as found in our dataset.


```
{
'AGE_GENDER':{
'__typename': 'WAISTUIAgeGenderType',
'__isIWAISTUIType': 'WAISTUIAgeGenderType',
'id': 'V0FJU1RVSUFnZUdlbmRlclR5cGU6NjUz',
'waist_ui_type': 'AGE_GENDER',
'serialized_data': '{
"age_min":6,
"age_max":53,
"gender":null
}',
'age_min': 6,
'age_max': 53,
'birthday': '***REDACTED***',
'gender': 'ANY',
'disable_preferences_entry': False,
'update_profile_uri': 'https://www.facebook.com/me/about/'
},
'LOCATION': {
'__typename': 'WAISTUILocationType',
'__isIWAISTUIType': 'WAISTUILocationType',
'id': 'V0FJU1RVSUxvY2F0aW9uVHlwZTpjb3VudHJ5LmhvbWUuREU=',
'waist_ui_type': 'LOCATION',
'serialized_data':'{
"location_granularity":"country",
"location_geo_type":"home",
"location_code":"DE"
}',
'location_name': 'Deutschland'
}
}
```


A.2 Examples of "Why am I seeing this ad?" across platforms

Figures 1, 3, 4 and 5 show screenshots of WAIST notices on Instagram, Facebook, X, and YouTube.

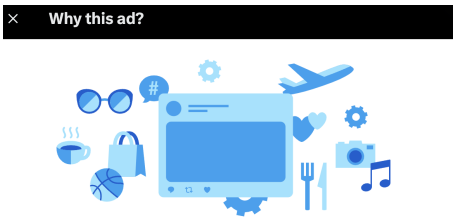

Figure 4: Screenshot of WAIST notice on X.

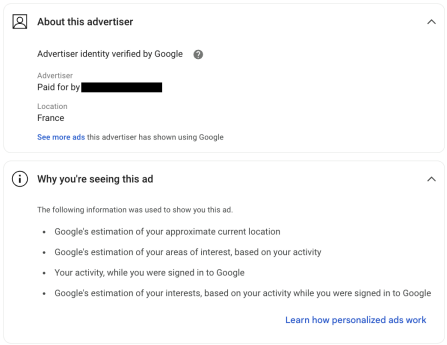

Figure 5: Screenshot of WAIST notice on YouTube.

B Browser Extension (Who Targets Me)

Who Targets Me (WTM) [28] is a browser extension for data collection through crowd-sourcing developed for auditing online advertisements—in particular, political ads—shown to users across major social media platforms (Facebook, Instagram, YouTube, and X). It works by detecting relevant HTML elements or intercepting network requests of the ads shown to users, then simulating user actions (such as clicking on information icons "i") to retrieve the *"Why am I seeing this ad?"* (WAIST) notices. The extension captures the data in the notices—including advertisement contents, identifiers, advertiser information, and, importantly, text explanations detailing the specific targeting parameters (e.g., demographic segments, inferred interests, or geographic boundaries)—and parses it into structured fields for analysis.

Table 6: Ad Explanation text examples per category for Instagram, YouTube and X. "advertiser" is a stand-in for the name of the entity promoting the ad.

Targeting category	Platform	Examples of ad explanations
Contextual	Instagram	/
	Youtube	The video you are watching
	X/Twitter	displayed near tweets containing the keyword apple
Demographic	Instagram	People who have indicated Bachelor's as their degree on Facebook.
	Youtube	Your gender / Your age / Home ownership status / Parental status.
	X/Twitter	The advertiser wants to reach people above the age of 21
Time/Location	Instagram	Connection information such as the IP address of your device
	YouTube	Your approximate current location according to Google's estimation
	X/Twitter	*advertiser* wants to reach people who are located here: Germany
Behavioral (Interests)	Instagram	Activities as part of Instagram/Meta usage, such as location tagging for photos, videos, or posts.
	YouTube	Google's estimate of your interests, based on your activity while signed in to Google.
	X/Twitter	*advertiser* wants to reach people interested in government
Remarketing	Instagram	People who have interacted with the Instagram account of *advertiser*.
	YouTube	Your use of the advertiser's website or app
	X/Twitter	You are a follower of *advertiser*.
Custom audience	Instagram	The advertiser uploaded a hashed list with your information to Facebook. We matched it with you.
	YouTube	The advertiser's goal to reach new customers who haven't bought from them before.
	X/Twitter	you belong to or resemble a target group of *advertiser*.
Lookalike	Instagram	People who resemble existing customers.
	YouTube	Your similarity to demographic groups that the advertiser is trying to reach – determined based on the activities recorded while you were signed in to Google
	X/Twitter	You resemble the followers of *advertiser*.

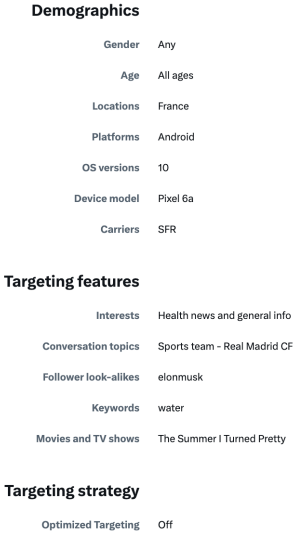

Figure 6: Screenshot of ad campaign summaries on X's Ad Manager, as presented to advertisers.

The data we acquired for this study is anonymized and does not contain any Personally Identifiable Information (PII).

C Examples of Campaign Summaries

Figure 6 shows a screenshot for a campaign summary as presented to advertisers on X. The information provided in such summaries could be directly communicated in ad repositories as "advertiser selection criteria".

D Facebook WAIST 2017-2025

Table 7: Evolution of targeting categories with WAIST attributes in Facebook ad explanations from September 2017 to March 2025.

Category	Fb attributes	Description	Time period	Phase(s)
Demographic	AGE_GENDER	Age ranges and gender	2017-2025	Pre-GDPR/CCPA, Post-GDPR & Post-DSA
	LOCALE	Local language settings	2017-2025	Pre-GDPR/CCPA, Post-GDPR & Post-DSA
	RELATIONSHIP_STATUS	Relationship information	2019-2025	Post-GDPR & Post-DSA
	ED_STATUS	Education level	2019-2025	Post-GDPR & Post-DSA
	WORK_JOB_TITLES	User occupation information	2019-2025	Post-GDPR & Post-DSA
	EDU_SCHOOLS	Education school	2019-2025	Post-GDPR & Post-DSA
	WORK_EMPLOYERS	Employer name	2019-2025	Post-GDPR & Post-DSA
Location	LOCATION	Home and temporary location	2017-2022	Pre-GDPR/CCPA, Post-GDPR
			2022-2025	Post-DSA
Behavioral	INTERESTS	user's interests	2017-2022	Pre-GDPR/CCPA & Post-GDPR
			2023-2025	Post-DSA
	BCT	Descriptors of a user action	2019-2025	Post-GDPR & Post-DSA
	FRIENDS_OF_CONNECTION	User's friends likes	2019-2022	Post-GDPR
Retargeting	CUSTOM_AUDIENCES_WEBSITE	Advertiser's website visit	2017-2025	Pre-GDPR/CCPA, Post-GDPR & Post-DSA
	CUSTOM_AUDIENCES_MOBILE_APP	Advertiser's mobile app use	2019-2025	Post-GDPR & Post-DSA
	CUSTOM_AUDIENCES_ENGAGEMENT_PAGE	Advertiser's Fb page interaction	2019-2025	Post-GDPR & Post-DSA
	CUSTOM_AUDIENCES_ENGAGEMENT_IG	Advertiser's IG interaction	2019-2025	Post-GDPR & Post-DSA
	CUSTOM_AUDIENCES_ENGAGEMENT_VIDEO	Advertiser's Fb video watching	2019-2025	Post-GDPR & Post-DSA
	CONNECTION	Advertiser's Fb page like	2019-2022	Post-GDPR
Custom Audience	CUSTOM_AUDIENCES_DATAFILE	Advertiser's uploaded hashed list	2017-2025	Pre-GDPR/CCPA, Post-GDPR & Post-DSA
Lookalike Audience	CUSTOM_AUDIENCES_LOOKALIKE	Lookalike custom audience	2017-2025	Pre-GDPR/CCPA, Post-GDPR & Post-DSA