

Defining Contextual Advertising

August 2024

Authored by

[Nathalie Maréchal](#), *Co-Director, Privacy & Data Project*

[Nick Doty](#), *Senior Technologist*

Over the past quarter-century, behaviorally targeted advertising¹ has become entrenched as one of the dominant business models for online publishers, social media platforms, mobile apps, and other web-based content, products, and services. A vast industry dedicated to collecting, analyzing, and monetizing data about virtually every aspect of people’s lives, comprising data brokers and various kinds of ad-tech intermediaries, forms the backbone of this ecosystem. The global digital advertising market is projected to reach US\$740.3bn in 2024.²

In response to a range of harms associated with behaviorally-targeted advertising,³ policymakers in the United States, European Union, and beyond have signaled an interest in using law and regulation to incentivize a shift toward more privacy-respecting advertising. In CDT’s view, policymakers should foster a competitive online advertising ecosystem that respects privacy and other human rights, supports independent media, enables the creation and availability of content—including from journalists and artists—and is accountable to the public, advertisers, publishers, and government regulators. Online advertising can and should be both privacy-respecting and economically viable.

Advertising in such an ecosystem could take many forms, one of which is contextual advertising. The initial discussion draft of the American Privacy Rights Act (the U.S. Congress’s latest attempt at bipartisan comprehensive federal privacy legislation) described contextual advertising as “when an advertisement is displayed online based on the content of the webpage or online service on which the advertisement appears.” But that neither represents a consensus understanding of contextual advertising nor answers a number of important questions about its scope. For example, some companies offer so-called “contextual advertising” services that rely on both contextual and behavioral data, as well as “cookieless” ad products that purport to

¹ For purposes of this document, we use the term “behaviorally targeted advertising” to mean the practice of serving users advertisements based on interests, characteristics and contexts inferred from their personal data.

² <https://www.statista.com/outlook/dmo/digital-advertising/worldwide>

³ A full accounting of the harms, benefits, and economic impact of various forms of behaviorally-targeted advertising is beyond the scope of this report.

respect individual privacy while replacing cookies with alternative persistent identifiers. CDT's stakeholder consultations in preparation for this report underscored the divergent views held by civil society and industry actors with respect to how "contextual advertising" should be defined.

To have a productive debate, and for law and regulation to establish understandable ground rules, we need clear, consistent, technically-grounded definitions that can be written into law and standards. While jurisdictions may ultimately vary in the restrictions they impose on different types of ads, public debate and policy-making will be well-served by common definitions. This report intends to fill that gap with respect to terminology about contextual advertising. CDT hopes that defining "contextual advertising" in dialogue with experts from both industry and civil society will be a first step toward developing a common understanding.

Many of the experts CDT interviewed for this report began by describing contextual advertising in opposition to what it is *not*: advertising that relies on the collection and tracking of personal and private information. Several people also described it by analogy to 20th century advertising paradigms that predate the internet: placing ads for luxury watches in a print magazine known for a high-income readership, or inserting ads for breakfast cereal between the Saturday morning cartoons watched by generations of children. Digital equivalents would be running an ad for sneakers alongside a news story about sports, or promoting a sale on snow shovels on a local news site in a region expecting heavy snow.

This is a sound starting point, but leaves important questions unanswered. What degree of geographic targeting should be permitted, if any, and how should a user's location be determined at the technical level? Can contextual ads be tailored according to the language a user is most likely to speak, and if so, what signals can be used for that prediction? How does the concept of "context," which most clearly evokes static web pages, translate to podcasting, streaming video, or social networking sites whose "feeds" vary for different people and at different times? This report, informed by extensive stakeholder consultations with experts from civil society and industry based in Europe and North America, proposes a definition for "contextual advertising" that answers these questions and, we believe, sets the stage for a more productive debate about the future of online advertising.

Proposed definition

ONLINE CONTEXTUAL ADVERTISING. – The term “online contextual advertising” means an advertisement that does not vary based on the identity of the intended recipient(s), and that is selected or targeted solely on the basis of one or more of the following factors:

- The immediate content of the user’s current webpage, app, video/audio programming, or online service on or within which the advertisement appears
- The immediate content of the user’s current request for information if the advertisement is displayed in proximity to the results of such search query or other request for information

Additional generally acceptable targeting criteria are:

- Such technical specifications, such as network or device information, that are necessary for the ad to be delivered and to display properly on a given device;
- The individual’s immediate presence in a physical location either 25 square kilometers or 10 square miles in size, or an area reasonably estimated to include online activity from at least 5,000 users, but no smaller than either 1 square kilometer or the area of a circle with a radius of 1,850 feet, including information about the current date, time, weather, and other current information about the location;
- The individual’s language preferences, as inferred from context, browser settings, other user settings, or geolocation.

Any data obtained for targeting a contextual advertisement should not be used to make further inferences about the consumer, to profile the consumer, or for any other secondary purpose. Consumers must be able to use technical means to obfuscate or change their physical location and to specify a language preference.

What is Context?

On the web, context is what is on the screen when someone visits a specific webpage (URL). Context can also include a search query, request for information or action being conducted, and the results for that query.

On a mobile app, context includes the app someone is currently and actively using, or a section thereof, but not details about the other apps installed on the device. Context on a mobile device is otherwise similar to on the web: the content someone is viewing, activity they're engaging in, a page or URL where the mobile app is showing web content, and any search queries or requests for information.

On a video or audio streaming service, context is the content that someone is currently watching or listening to, including content streamed immediately before the ad or that is about to play. Streaming context is limited to details about the channel and exact show someone is consuming; "watch history" (beyond the piece of content that played immediately prior to the ad) should not inform the placement of contextual ads.

On a social media platform, context is the specific URL within the platform, such as a page or community, and the content that appears on that URL. Context can also include a search query, request for information or action being conducted within the platform, and the results thereof, as well as audio or video content streamed immediately prior to the ad or that is about to play. Large amounts of personal information, as well as additional context about what content is being shown in someone's feed, are typically used to target advertisements and sponsored content on social media. Social media is the most difficult environment for defining contextual advertising, in no small part because contextual ads are somewhat antithetical to social media platforms' predominant data-intensive business models. CDT nonetheless calls for context on social media platforms to exclude declared or inferred personal data not contained within the platform URL, including users' profile information such as gender or age.

Our research identified three main methods for contextual targeting in a social media environment:

- 1) Placing an ad within a thematically relevant portion of the platform, such as a community, group, or page devoted to a particular subject;
- 2) Showing an ad about a topic next to (adjacent on screen, or immediately before, after or during) a piece of content relevant to that topic;
- 3) Selecting an ad in response to a search query, request for information or action being conducted within the social media platform, and the results thereof.

With respect to temporality, and regardless of medium, context is limited to the media content that someone is consuming in a specific moment: what is on the screen when someone visits a given webpage (URL), a specific search query and the results for that query, the app someone is currently and actively using (or a section thereof), or audio or video content streamed

immediately prior to the ad (“post-roll ad”) or that the user has selected to play next (“pre-roll ad”). While browsing, viewing and search histories’ may be useful for ad targeting, we maintain that compiling and processing such activity logs constitutes behaviorally targeted advertising, not contextual advertising.

Generally Acceptable Targeting Criteria

Using certain types of data should be permitted either for technical functionality or because they support matching audiences with relevant ads without presenting significant privacy risks. CDT understands these types of data as part of the typical context of an online interaction: device information necessary for rendering, coarse (non-precise) location of delivery, and language selection.

Technical functionality

At a high level, adtech systems use certain technical details about a user’s device to deliver and display ads properly, a process known as “content negotiation.” Information provided for content negotiation is appropriately used to select advertising in a format that the user’s device can most easily access, which may include file format preferences, device capabilities, or even preferences around visual display or network usage, but not unique device identifiers such as IMSI and IMEI numbers.

However, such device data can be also used for device fingerprinting (e.g. using a device’s combination of technical details and settings as a unique identifier) or to support discriminatory ad-targeting, for example offering different products or price points to users of more expensive devices. Using content negotiation information to re-identify a user or to target advertising based on that user’s past activity is plainly not contextual advertising. Using explicitly provided content negotiation information to make inferences about a user for ad targeting is also not contextual advertising, except for the specific categories noted below: geolocation, context derived from geolocation, and language.

Geolocation

Many of the stakeholders we consulted agreed that a practicable definition of contextual advertising should allow for the use of some form of “coarse” (non-precise) geolocation. At a minimum, such geolocation may be necessary for legal compliance in certain circumstances. There was also broad agreement that small businesses and political campaigns in particular benefit from the ability to reach audiences in a specific geographic area, and that consumers can also benefit from seeing ads that are relevant to their part of the world.

On the other hand, civil society experts emphasized that precise geolocation data is among the most sensitive forms of data that can be collected, especially when companies indefinitely retain records of a person or device’s movements over time. Many people can be uniquely identified based on a few locations where they spend the majority of their time, such as the

location of their home and work. Location data initially collected to provide either a service requested by a user or for advertising purposes can be sold to data brokers and ultimately used by other actors such as commercial spyware vendors with harmful real-world impact. It can be obtained by government authorities either pursuant to a bona fide legal process, by purchasing it from a location-data broker, or by simply asking. This is of heightened concern in the U.S. in the post-*Roe* era.

There was less agreement on where the line between permitted coarse geolocation and prohibited precise geolocation should be drawn, however. The Banning Surveillance Advertising Act, first introduced in the U.S. Congress in 2022, would allow contextual ads “to be targeted to an individual, connected device, or group of individuals or connected devices based on a recognized place associated with the individual, connected device, or group of individuals or connected devices.”⁴ Several U.S. states and the initial discussion draft of the 2024 American Privacy Rights Act define precise geolocation data as “information that reveals the past or present physical location of an individual or device with sufficient precision to identify (A) street-level location information of such individual or device; or (B) the location of such individual or device within a range of 1,850 feet or less.” Industry representatives interviewed by CDT generally endorsed this definition, which some civil society groups viewed as still too precise, instead proposing “a physical location no smaller than 10 square miles.”

Civil society and industry stakeholders alike noted that any area-based limits on geotargeting should account for differences in population density between rural and urban areas, with more precise geotargeting being permitted in more populated areas. Recognizing that such line-drawing is, by necessity, somewhat arbitrary, and that jurisdictions that use different systems of measurement (i.e., metric vs. U.S. customary units) will prefer to specify thresholds expressed in whole numbers, CDT recommends the following as thresholds for “coarse” geolocation information that may be used in connection with contextual advertising:

- In general, geotargeting may be as precise as an area of 25 square kilometers or 10 square miles.
- Geotargeting for densely populated areas may be more precise, if the geotargeting area includes online activity from at least 5,000 users and does not constitute precise geolocation, which is defined as “a geographic area that is equal to or less than the area of a circle with a radius of 1,850 feet” (roughly equivalent to 1 square kilometer) under state law in California and some other U.S. states. Cities with such a population density include parts of New York City (Manhattan), Miami, Paris and Barcelona.
 - This does not mean that 5,000 people will receive the ad in question, but that 5,000 people would be eligible to be shown the ad on the basis of their current physical location. A subset of that population would see the ad, depending on what they are viewing or doing online, the specific contextual targeting criteria, and the outcome of any ad auction.

4

<https://www.congress.gov/bill/118th-congress/house-bill/5534/text?s=7&r=1&q=%7B%22search%22%3A%22Banning+Surveillance+Advertising+Act+of+2023%22%7D>

Beyond drawing the line between permissible and sensitive geolocation, the question remains of how a given device's location should be determined. Today, geolocation is most often determined based on a device's Internet Protocol (IP) address, which often maps roughly onto a geographic area even though it more accurately reflects digital routing infrastructure than geographic areas. For mobile devices, both GPS and triangulation based on cell phone towers can also be used. However, IP addresses can be used to identify and track individuals and/or households. One civil society group recommended the creation of new technologies, perhaps at the browser level, that would transform IP-derived location information into a city/regional level. Rural areas would need this area to be larger than city areas to prevent identification as well as discriminatory targeting. Time-zone information—an incomplete version of geolocation—can also be derived from a device's time-clock and used to tailor ads based on the time of day.⁵

Several civil society respondents expressed a strong view that geolocation should only be determined on the basis of self-declared information. While CDT finds this somewhat impractical in the short term, we do agree that users should be able to use technical means, such as a Virtual Private Network (VPN), to hide or change their geolocation. If such means are used, self-declared geolocation should override actual geolocation for purposes of targeting contextual ads.

Context derived from geolocation

Permissible geolocation can also be used to derive additional information about a user's context that, in CDT's view, should be allowed to inform ad targeting. For example, geolocation is tied to weather, time of day, and sociocultural events like widely observed holidays and other events. Retailers should be able to run "Back to School" ads in different cities based on when the local school district's school year starts, for example. Similarly, a sandwich chain could place ads for bagels in the morning and burgers at lunchtime. CDT's position is that this kind of tailoring should be permitted within the definition of contextual advertising.

Language

Common sense suggests that most advertisers would prefer to show ads in a language that their audiences will understand, and that most people would also prefer to receive communications they understand. This means that adtech systems may need to make inferences about the language(s) that ad recipients will prefer. Today, this is overwhelmingly done on the basis of geolocation: the same device will receive ads in French when its IP address suggests it is located in Paris, and in English when it (and its owner) visits Washington, DC. Building on the earlier point about geolocation and legal compliance, some countries require advertisements to use their national language (France is one such example).

⁵ Implementation details vary for smartphone operating systems and the web. Methods have changed slightly over the years but it's still accessible to developers for a variety of reasons.

Language preferences can also be inferred from contextual data—while reading a news article written in Chinese, the user sees ads in Chinese. Finally, language can be derived from user settings at the device, browser, application, or website level. For example, if a user sets a social networking site’s interface language to German, that person may be more likely to see ads in German. Adtech systems should be able to infer language preferences from any of these three sources (geolocation, context, and user settings). But as with geolocation, users should be able to use technical means to communicate their advertising language preferences, which should override any other signals.

Other Considerations

When the context itself may be “sensitive data”

Some U.S. state laws, notably Washington’s “My Health, My Data Act,” prohibit including the fact that someone visited a health-related website in that person’s adtech profile.⁶ CDT supports that protection. However, we believe that it should be permissible to place contextual ads related to the specific topic on a health-related webpage or website, as long as no personal information is tracked or used. For example, a manufacturer of insulin pumps should be able to place ads on a webpage offering advice about diabetes management. The same should be true for other contexts that may involve sensitive content. Contextual advertising is the least privacy-invasive way for the manufacturer to reach relevant audiences and for the publisher to generate advertising revenue, both legitimate interests that online advertising supports.

That being said, publishers should conduct due diligence to avoid running ads that would be particularly harmful to their publication’s audience, and adtech intermediaries should provide publishers with the tools to do so. For example, a website dedicated to sobriety should not display ads for alcohol. This issue is better addressed through industry self-regulation and through internet standards than through law, however.

Compliance with laws protecting minors

Several industry respondents highlighted that in some jurisdictions, they need to “know” whether a user is a minor to comply with content or targeting restrictions aimed at protecting children. Depending on the knowledge standard in the statute in question, advertisers and their service providers may be able to comply with the law by placing their contextual ads in appropriate environments and/or refraining from using personal data, as is required by the EU’s Digital Services Act. In any case, age should not be used as a targeting criterion for contextual ads.

Use of data beyond targeting

While this report focuses on contextual *targeting* of advertising, experts have also noted that personal data is often collected or used for other advertising-related purposes, including legal

⁶ <https://iapp.org/resources/article/washington-my-health-my-data-act-overview/>

compliance, fraud mitigation, frequency capping, and measurement (including attribution), even when the targeting is done solely on the basis of contextual factors, as in our proposed definition above. Recommendations for what limits should be placed on the collection, retention and use of personal data for such purposes are beyond the scope of this report. In many cases, it may be reasonable for data to be used for such purposes to the extent that it is necessary, limited and proportionate to those purposes; however, where reasonable systems exist to provide aggregated privacy-preserving methods to accomplish the general purpose, it is not reasonable to continue use of individuals' data.

Consent and self-directed advertising

Finally, several experts raised the issue of consent: can internet users consent to having additional data used to tailor the contextual ads they see? CDT supports permitting self-directed advertising, where users themselves volunteer their demographic information to inform the ads they see, as long as providing this information is truly optional and can be revoked at any time, in compliance with relevant laws such as GDPR. However, self-directed advertising is distinct from contextual advertising.

**Find out more about CDT's work as part of the
Future of Online Advertising Project at
cdt.org/online-advertising.**

Acknowledgements

CDT is grateful for the thoughtful input of numerous stakeholders from industry and civil society who are members of the working group formed as part of our [Future of Online Advertising Project](#).

The Center for Democracy & Technology ([CDT](#)) is the leading nonpartisan, nonprofit organization fighting to advance civil rights and civil liberties in the digital age. We shape technology policy, governance, and design with a focus on equity and democratic values. Established in 1994, CDT has been a trusted advocate for digital rights since the earliest days of the internet.