

Topics API: Criteo's First Look at Google's Interest-Based Advertising Solution

Google [announced](#) at the end of July that they intend to begin phasing out third-party cookies for their Chrome browser in the second half of 2024. As a potential solution, Google [launched](#) an initiative in 2019 called [Privacy Sandbox for the web](#).

Since the very beginning of this initiative, [Criteo](#) has been an active participant in various testing phases of its Privacy Sandbox, including the first "Origin Trial" in 2021. During this first trial, we shared [our feedback on the FLoC API](#), which was eventually discarded by Google, on July 13, 2021.

In 2022, Google launched a new series of Origin Trials across various Privacy Sandbox proposals. As a leading provider of performance-based ad technology solutions for media owners and marketers, Criteo is committed to testing the proposals that might support open web advertising core functions including:

- audience building and targeting, with FLEDGE and Topics APIs;
- attribution and measurement, with the Attribution Reporting API.

Our objective is to test Google's Privacy Sandbox to measure and maximize the value delivered to advertisers and publishers. We will quantify the utility of the APIs in comparison to the status quo solutions powered by third-party cookies. For each API we will rigorously study the foreseeable impacts on publishers and advertisers, share factual results, and provide recommendations to help improve the proposals and shape the success of the Privacy Sandbox.

One of the key solutions being proposed by Google is the **Topics API**, which seeks to enable interest-based targeting for advertisements on the open web once Google ends support for third-party cookies. The [Topics documentation](#) defines interest-based advertising in the following way:

“Interest-based advertising (IBA) is a form of personalized advertising in which an ad is selected for the user based on interests derived from the sites that they’ve visited in the past. This is different from contextual advertising, which is based solely on the interests derived from the current site being viewed (and advertised on). One of IBA’s benefits is that it allows sites that are useful to the user, but perhaps could not be easily monetized via contextual advertising, to display more relevant ads to the user than they otherwise could, helping to fund the sites that the user visits.”

This article will outline our testing process around the relevancy of Topics. We were specifically interested in assessing how Topics can enable advertisers to create interest-based audiences as they do today with third-party cookies. It’s important to mention that this is the first version of the Topics API and there will be room for many iterations and improvements. Our objective is to share Criteo’s early findings on the first iteration of Topics to help improve the proposal.

TL;DR

Criteo implemented Google’s Topics API for interest-based advertising on the open web without third-party cookies. We began testing in June with Google opening up traffic to 1% of Chrome’s stable version on desktops in August, and mobile Android in September. We are now observing close to 1 million daily Topic-tagged users and can offer our first assessment of the value of this signal in the present article.

After a few months of data collection regarding the Topics API, we identified some limitations in the solution’s ability to emulate interest-based audiences without the use of third-party cookies. **Overall, we observe that Criteo’s interest-based audiences are five times more relevant than those generated by this first iteration of the Topics API.** While Criteo will continue to perform experiments to evaluate how this would impact marketers (ROI) and publishers (Revenue), a negative impact can be anticipated for both considering the drop in audience relevancy.

We believe the following improvements if put in place could increase the relevancy and utility of Topics for interest-based advertising:

- Use a more detailed and comprehensive taxonomy with a special focus on commercially relevant topics

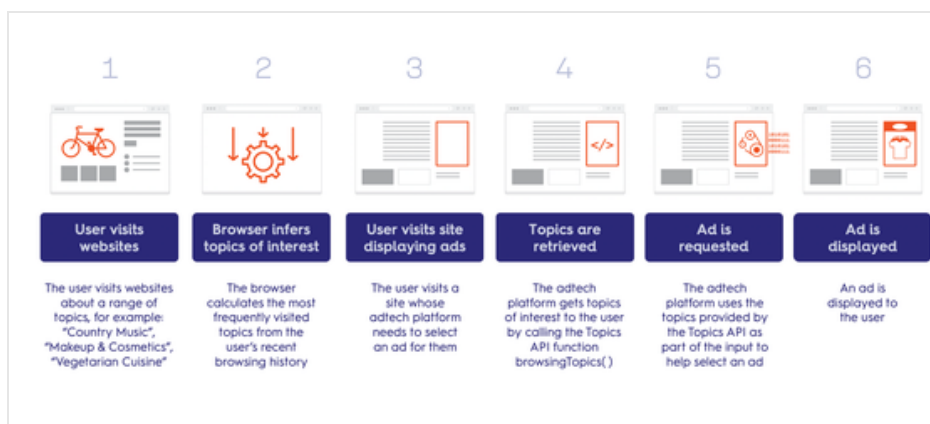
- Identify the most frequent topics per user in comparison to the general population
- Classify the web more granularly by taking into account URL and Page content
- Shorten the epoch duration to emphasize recently viewed topics

We are also aware that the signals the Topics API provides could have further utility if combined with other signals such as first-party data and contextual information, and this will be a matter for our future experimentation.

Criteo’s commitment has been and will continue to be to test and provide constructive feedback so that Google, the industry, marketers, and media owners can act and adapt positively to this upcoming change.

The Topics API

The Topics API is meant to power interest-based advertising on the open web by inferring a user’s top interests (or topics) and making those topics available to any marketer. To enable this, Google classifies websites into one or more topics from this [taxonomy](#) using a mix of manual and automated classification. Then, each week (epoch), and for a total of three weeks, Google will retrieve a user’s five most frequent topics based on his/her browsing history. Lastly, the advertising community of players (publishers, advertisers, DSPs, SSPs, etc.) calling the Topics API on a given website may retrieve, at most, one of the topics per epoch, up to a maximum of three topics. These retrieved topics can then be used for interest-based advertising.



For more details on how Topics API works, check out [Google's blog post on the Topics API](#).


What Criteo observed in initial testing

Shortly after Google announced the deployment of the Topics Origin Trial, Criteo reacted with a hands-on approach to implement and start testing. We began logging the first events that contained users' topics in June and saw significant volumes as Google opened up traffic to 1% of Chrome's stable version in August.


Figure 2 illustrates the evolution of the number of Topics-enabled users in the last month. We see that traffic ramped up at the beginning of September and, **as of today, we observe more than 1 million distinct daily users per site containing at least one topic tag, which was enough to draft our first conclusions.**



After the number of topics volume increased, we could already visualize the most common topics observed by the advertisers and publishers we are integrated with. Our first observation, illustrated in Figure 3, shows that more than 60% of observed topics belong to Level 1, the most general category within the taxonomy, containing only 24 topics in total.



Furthermore, the three main Level 1 topics observed are “Arts & Entertainment”, “News” and “Shopping”. If we include “Online Communities”, which is the 9th most observed level 1 topic, **we observe that almost half of all topics observed are not tied to commercially valuable information about what might interest a user’s shopping behavior.**



This phenomenon is a direct consequence of three important characteristics of the Topics API:

First, the **current taxonomy, which contains 349 topics, is missing some very relevant commerce-related topics such as shoes and dresses.** Fortunately, Google has already shared with us that this is only a testing taxonomy and is planning to enhance and enrich it in the coming months, which we will closely monitor. As Google makes a more granular taxonomy available Criteo will re-test the solution to understand the impact.

Second, the **mechanism** of how topics are assigned to users rarely observes niche topics. As mentioned earlier, Google selects the top 5 topics a user has browsed during the past 7 days using absolute frequency. We found that this approach tends to favor more popular and broad categories of interest such as “Shopping”, “Online Communities” (social media), or “News”. In contrast, more niche topics such as “Car Rentals,” which are more definitive signals usually

associated with immediate commerce intents and thus more commercially valuable for marketing purposes, are rarely observed.

Finally, the third reason is the Topics **classifier** itself. We have taken a closer look into this below.

Topics classifier model

The importance of URL and page content signals

To assess the Topics classifier, we compared it with Criteo’s classifier, which is mostly used to predict the theme and sentiment of publisher web page contents. Our classifier uses natural language processing techniques to interpret a URL and, when allowed, a crawler extracts the main content of the page.

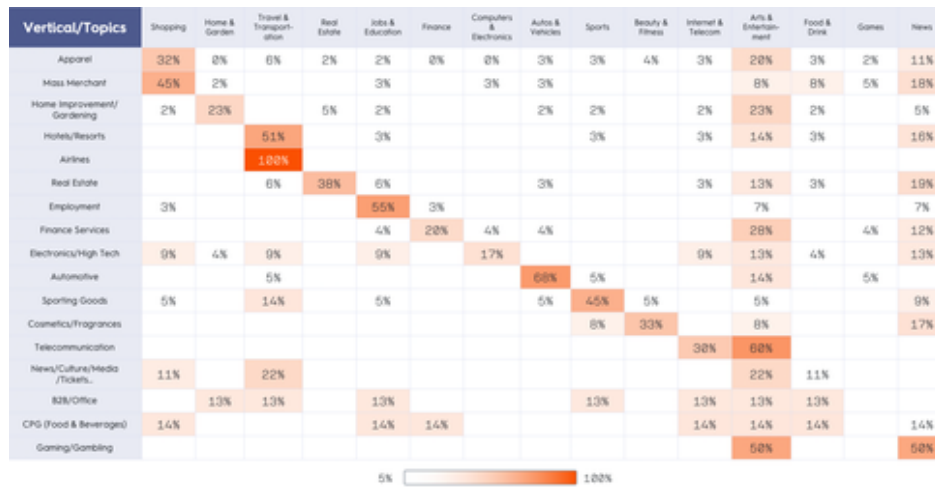
Advertiser or publisher	Sub-category	Criteo’s classification	Topics classification
Amazon	Gaming	Video Gaming Console Games	Shopping
Ebay	Home & Garden	Home & Garden	Shopping
Fox News	Sports	News & Politics Sports	News & Politics
MSN	Finance	Personal Finance Personal Investing Stocks & Bonds	Arts & Entertainment Internet & Telecom News

One of the limitations of Google’s Topics classifier is that it only considers the hostname (including domain and subdomain) of the website for its classification. What this means is, even if the user is browsing a sports section within a mass merchant retailer, the classifier will only consider the hostname, and will therefore output “Shopping” as the main topic for that interaction. Table 1 illustrates this situation using concrete examples of advertisers' and publishers' websites. This is a relevant point for interest-based advertising since much of a user’s online activity occurs in sub-categories within websites of large general retailers, publishers, or

social networks. Consequently, there are a lot of signals about the user’s interests that Topics is not capturing.

Advertiser classification

In order to assess the quality of the Topics classifier, we compared it with Criteo’s advertiser vertical, a human-assigned category for each advertiser. To avoid taking into account the limitation generated by the testing taxonomy, we only considered the advertisers' verticals that have one or more equivalent (or similar) level 1 topic(s). Figure 5 displays a heatmap that contains the classification of topics corresponding to each vertical for Criteo’s top 1000 advertisers.



With the exceptions of the “B2B / Office”, “CPG (Food & Beverages)” and “Gaming / Gambling” verticals, we observe that the dominant topics for each vertical make sense.

However, it does stand out that many advertiser verticals tend to be significantly classified within the “Arts & Entertainment” and “News” topics.

Evidently, the predominance of these topics, as shown in Figure 4, could be explained partly by the fact that these topics are largely browsed by users, and partly by the false positives shown in Figure 5.

Publisher classification

Next, we compared Criteo's publisher classifier and the Topics classifier on a set of well-known publishers.

Advertiser or publisher	Criteo's classification	Topics classification
yahoo.com	News & Politics Personal Finance Personal Investing	Arts & Entertainment Internet & Telecom Search Engines News
weather.com	News & Politics Weather	Weather
msn.com	Technology & Computing Computing Internet	Arts & Entertainment Internet & Telecom News
newsroomlabs.com	News & Politics	Science News
microsoftcasualgames.com	Technology & Computing Video Gaming	Arts & Entertainment Computer & Video Games
tenki.jp	News & Politics Weather	Weather
foxnews.com	News & Politics	Arts & Entertainment News

We observed that Topics performs relatively similarly to Criteo's classifier on publisher sites despite having a less sophisticated taxonomy.

Criteo's recommendation for the next stage of the Topics classifier

Based on our analysis, the Topics classifier consistently assigns the same topic to each hostname. However, because the Topics classifier only takes into account the website's hostname, it creates a disadvantage to the model's utility. Including both, the **complete URL**, and the **page content**, would generate a significant improvement in this regard.

Assessing the relevancy of Topics Audiences

Now, we get down to business. Here, we explore how relevant topics-based audiences would be to generate effective results for marketers compared to existing solutions for interest-based advertising. It is worth stating that the deprecation of third-party cookies would most certainly have an impact on the predictive bidding models used across programmatic advertising. In this section, we will focus exclusively on how the Topics API affects the industry’s upper-funnel audience-building capabilities.

Currently, Criteo, like many other DSPs, creates interest-based audiences powered by third-party cookies, which we call In-Market Audiences. Our marketer buying platform relies on multiple features to calculate the user’s commerce interest. This allows advertisers to choose among different interest and profile configurations to build a precise interest-based audience. For instance, a marketer may decide to target the “in-market for a specific shoe brand” audience. In addition, marketers can enhance these audiences by using machine learning models that deduce the user’s gender or purchasing power. Overall, the following configuration settings exist and can be mixed and matched:

- Commerce category
- Brand
- Purchasing Power
- Gender

The deprecation of third-party cookies will certainly make it challenging (if not impossible) to create most of these audience configurations, since Topics can only replicate, to a certain extent, In-Market audiences that are purely commerce category-based. Therefore, as we see in Table 3, **around 51% of publisher revenue is generated by audiences that are not reproducible with today’s version of Topics.**

Features used in audience configuration	Share of publisher revenue
Only Commerce category (partially replicable by Topics)	49%
Commerce category + other features	43%
Other features	8%
TOTAL	100%

It is fair to say that category is the most relevant of these configurations. In fact, **In-Market audiences with at least one category filter represent close to 92% of In-Market audience publisher revenue.** What this means is, although we cannot replicate most configurations, it will still be very beneficial to have the ability to reproduce category-based In-Market audiences using Topics-based audiences. And that is exactly what we at Criteo attempted to do!

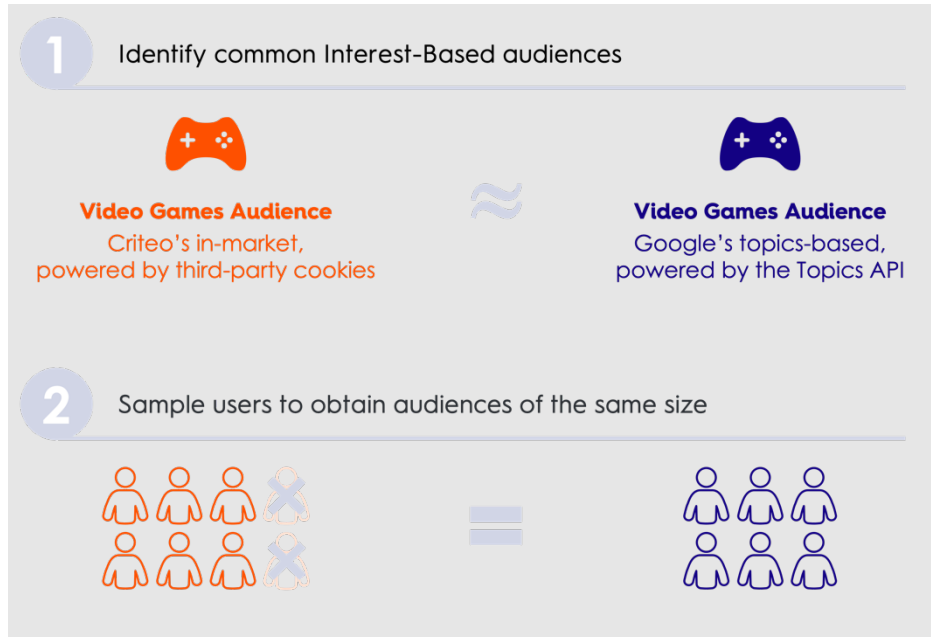
Criteo’s In-Market audiences vs. Topics-based audiences

Experiment Set-Up

Given the limited number of and lack of granularity for commerce-related topics in the current Topics taxonomy, we selected only Level 1 commerce categories (albeit, the most general) within our In-Market audiences configurations that have a matching or similar topic. These are detailed in the following list along with their share of publisher revenue.

Google’s category level 1 for Criteo’s in-market audiences	Closest topic(s) from Topics taxonomy [TOPIC taxonomy ids]	Share of publisher revenue
Apparel & Accessories	Shopping [289.299]	50,8%
Electronics	Computer & Electronics [124.146]	13,8%
Home & Garden	Home & Garden [207.215]	12,8%
Health & Beauty	Beauty & Fitness [86.100]	5,1%
Arts & Entertainment	Arts & Entertainment [1.57]	2,8%
Toys & Games	Games [180.195]	1,3%
Vehicles & Parts	Autos & Vehicles [57.87]	1,1%
Luggage & Bags	Travel & Transportation/Luggage & Travel Accessories [334]	0,8%
Business & Industrial	Business & Industrial [103.125]	0,7%
Food, Beverages & Tobacco	Food & Drinks [172.179]	0,5%
Animals & Pet Supplies	Pets & Animals [263.271]	0,2%

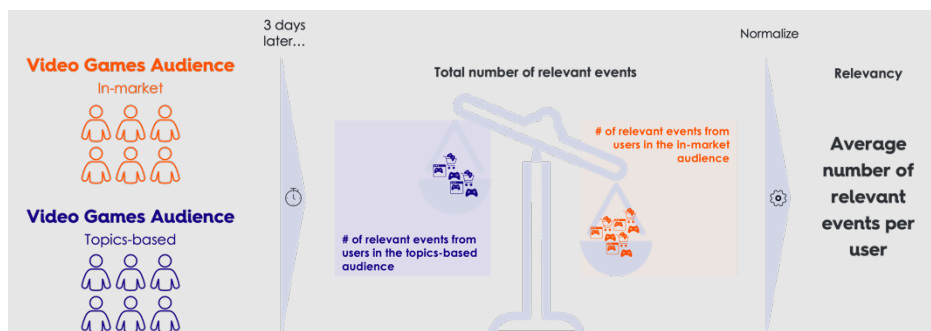
Our next step was to build the appropriate audiences. For Topics, we retrieved all those users that had the respective topic tag in the past 7 days, while for the In-Market audience we considered the most basic configuration, one that only considers users that browsed the given category in the past 7 days. No predictive techniques nor machine learning were leveraged in our benchmark audience. To make these audiences comparable, we randomly sampled the In-Market audience to match the number of users within the equivalent Topics-based audience.



Audience Relevancy


When an advertiser configures a campaign to deliver to an audience interested in a given category, they are looking to engage people who will likely continue to browse, be interested in, or purchase within this category in the following days.

Therefore, we define the **relevancy** of an audience as the number of **relevant events** observed posterior to the audience creation. Relevant events are organic visits, baskets, and purchase events associated with products of the category selected in the audience configuration. Figure 7 explains the relevancy calculation for a “Video Games” audience.



Results

In the following chart, we present the number of relevant events in the posterior 3-day period performed by users within Criteo's In-Market and Topics-based audiences.



We observe that Criteo's current solution for interest-based advertising generates audiences that are on average 5 times more relevant than those enabled by the current version of the Topics API. For instance, in Electronics, a popular In-Market category: we observe 270 relevant posterior events per 1000 users with the Topics audience, while for the In-Market one we record 830. This translates into a x3 relevancy uplift.

At this stage, we are not able to quantify what this difference in audience relevancy would generate for marketers and publishers, but it is certainly something we would like to get to in future experiments. Our experience in building audiences tells us that marketers are likely to experience a sizeable decrease in the desired commerce outcomes relative to today's third-party cookie-powered solutions. This could put publishers at risk as marketers may shift spending to environments not impacted by signal depreciation.

User Fatigue: "Stop Showing me Shoes!"

Having topics assigned only after one week, and with a duration of three weeks, is not adequate to capture timely interest buying intentions. Studies on the matter have widely shown that event

recency is very important in the predictive capacity of an audience. With the exception of expensive durable goods (e.g., refrigerators), people who are buying online tend to have a short consumer journey: it is unusual for someone to spend 3 weeks browsing for shoes, for example. This requires marketers to act quickly if they want to attract this buyer. Even if Topics perfectly captured commerce intentions, an advertisement display motivated by a 3-week-old browsing event will probably be anything but noise to the user.

Conclusion

After a few months of data collection regarding the Topics API, we identified some limitations in the solution's ability to emulate interest-based audiences without the use of third-party cookies. Since this is just the starting point of the Topics proposal, we expect and hope many improvements will come from utility testing and feedback from all testers. Until then, we are providing a series of improvements that we believe will have a positive impact on the next iteration of this API utility:

- **Use a more detailed and comprehensive taxonomy with a special focus on commercially relevant topics:** Generic Topics such as “Shopping” or “News” lack commercial value for advertisement purposes. The first testing taxonomy is limited and excludes many relevant commerce categories that are essential for the advertisement business. We suggest a more granular commerce taxonomy such as the [Google product taxonomy](#).
- **Identify the most frequent topics per user in comparison to the general population:** The process that defines how the top 5 topics are assigned to users tends to favor the same generic and popular topics for most users. We propose implementing a relative frequency that adjusts to the topic's overall popularity within the user's country or region.
- **Include URL and Page content in the Topics classifier features:** The classifier assigns descriptive categories to advertisers' and publishers' hostnames. As per Criteo experience in building audiences, we are confident that taking into account the URL and the content of the website as features in the classifier will allow for much more accurate classification of websites' pages and thus improve the signaling within Topics
- **Improve timeliness in topic selection.** As demonstrated, a week-old behavior has limited utility for an online advertisement. Therefore, we propose that topics be assigned to the user after the last 2 days of browsing history.

We are confident that, should these four suggestions be implemented, the relevancy of topics-based audiences would be considerably better.

Next Steps

As we have seen in this article, the Topics API on its own and in its current form likely does not adequately address marketers' needs for interest-based advertising. But it is possible that this solution, in combination with other signals, such as contextual features, first-party data, and those provided by FLEDGE and the Measurement APIs, could indeed provide a valuable alternative.

Criteo's commitment is to continually add complexity to our experiments and analyze the performance of combined solutions as soon as this is technically feasible. We welcome Google's openness to the constructive feedback we provided on this very first version of the Topics API. We will continue to test and provide recommendations so that Google, the industry, marketers, and media owners can act and adapt positively to this upcoming change.

Stay tuned for our next article!