

## A Definitions of Targeting Types

In this section, we provide the terminology and definitions for the 16 targeting types we investigated in Part 1 of our user study. We took all terminology and definitions that we showed participants verbatim from Twitter Business's help pages (<https://business.twitter.com/en/targeting.html> and pages linked from it).

- **Age targeting** allows advertisers to target people by age buckets, such as 18+ years old or 18-24 years old.
- **Behavior targeting** allows advertisers to target people based on inferred behavior, such as shopping and lifestyle habits or income.
- **Conversation topic targeting** allows advertisers to target people based on topics they have engaged with (e.g., Tweeted, clicked, Retweeted, replied, liked, viewed) on Twitter.
- **Event targeting** allows advertisers to target people based on events they are interested in or have engaged with (e.g., Tweeted, clicked, Retweeted, replied, liked, viewed) on Twitter.
- **Follower lookalike targeting** allows advertisers to target people who don't necessarily follow a given account, but have similar interests or demographics to the account's actual followers.
- **Gender targeting** allows advertisers to target people based on their self-reported or inferred gender.
- **Interest targeting** allows advertisers to target people based on inferred interests, as determined by who they follow on Twitter and their Tweets, Retweets, and clicks.
- **Keyword targeting** allows advertisers to target people based on words or phrases they have Tweeted about or searched for on Twitter.
- **Language targeting** allows advertisers to target people who use a certain language on Twitter.
- **Location targeting** allows advertisers to target people based on region, city, metro or zip code.
- **Mobile audience targeting** allows advertisers to target people who use their mobile app.
- **Movie and TV show targeting** allows advertisers to target people based on movies and TV shows they have watched or are likely to watch.
- **Platform targeting** allows advertisers to target people who use a certain platform, such as iOS or Desktop, to access Twitter.
- **Retargeting campaign engager targeting** allows advertisers to target people based on prior engagement with (e.g., Tweeting, clicking, Retweeting, replying, liking, or viewing) their company.
- **Tailored audience (list) targeting** allows advertisers to reach specific people on Twitter by uploading lists, which contain personal information (email addresses, phone numbers, or Twitter handles) that are matched to Twitter users' accounts.
- **Tailored audience (web) targeting** allows advertisers to target people who have visited their website.

## B Raw Data Example

```
[ {
  "ad" : {
    "adsUserData" : {
      "adImpressions" : {
        "impressions" : [ {
          "deviceInfo" : {
            "osType" : "Ios",
            "deviceId" : "#####",
            "deviceType" : "iPhone X"
          },
          "displayLocation" : "SearchTweets",
          "promotedTweetInfo" : {
            "tweetId" : "#####",
            "tweetText" : "RT @SpotifyBrands: Young people's digital lives are subtly shifting
                          culture. Discover more with our global trends report.",
            "urls" : [ ],
            "mediaUrls" : [ ]
          },
          "advertiserInfo" : {
            "advertiserName" : "Spotify",
            "screenName" : "@Spotify"
          },
          "matchedTargetingCriteria" : [ {
            "targetingType" : "Events",
            "targetingValue" : "Back to School 2019"
          }, {
            "targetingType" : "Age",
            "targetingValue" : "18 to 49"
          }, {
            "targetingType" : "Locations",
            "targetingValue" : "United States"
          }, {
            "targetingType" : "Platforms",
            "targetingValue" : "iOS"
          } ],
          "impressionTime" : "YYYY-MM-DD HH:MM:SS"
        } ]
      }
    }
  }, {
    "ad" : { ...
```

## C Requests For Clarification of the Data Files We Made to Twitter Under GDPR

In this section, we describe the communications between a member of our research team and Twitter through which we attempted to confirm and clarify the meaning of fields present in users' data files downloaded from Twitter. This individual invoked their GDPR rights as an EU citizen as the basis for these requests and communications.

### C.1 Initial Request

A member of the research team (an EU citizen) submitted the following request and series of follow-ups to both of the following channels:

- We sent the request below to Twitter Support via the Twitter Privacy Inquiries online form indicated in Twitter's Privacy Policy (<https://help.twitter.com/forms/privacy>). We sent this request on **June 26, 2019**.
- After we received an unsatisfactory response to the initial request (below, "Twitter's Initial Response") on July 7, 2019, we sent the same request to Twitter's Data Protection Officer via the online form indicated in Twitter's Privacy Policy ([https://twitter.ethicspointvp.com/custom/twitter/forms/data/form\\_data.asp](https://twitter.ethicspointvp.com/custom/twitter/forms/data/form_data.asp)). We sent this request on **July 9, 2019**.
- We received a satisfactory response from Twitter Office of Data Protection on **November 15, 2019**.

The member of the research team included the following data they downloaded about their own Twitter account: the PDF file containing the list of "Similar audiences" and "Tailored audiences," as well as the "ad-impressions.js" file.

Hello,

I have downloaded my Twitter data from the settings page. The information displayed in this exported data or in various help center articles does not fully explain about how my personal data is being used for advertising purposes, and I ask for a number of clarifications:

1. I have downloaded my data, and in the zip file there is a file called "ad-impressions.js", which I have also submitted with this request as a courtesy. For each ad in this file, there is a field called "matchedTargetingCriteria", and I would like to understand how to interpret this information.
  - (a) Does matchedTargetingCriteria represent the criteria that the advertiser chose to target this ad? If not, what does it represent?
  - (b) If (a), does the matchedTargetingCriteria shown for a single ad represent ALL the criteria that the advertiser chose for that ad, or are there criteria that the advertiser chose when targeting this ad that are not reflected in the export?
  - (c) Do all the matchedTargetingCriteria apply to me, e.g. if an ad is targeted to "Follower look-alikes: @Twitter, Keywords: Privacy, and Locations: United Kingdom", does that mean that I am seeing this ad because ALL of those criteria apply to me, or only a subset?
    - i. If it is only a subset, how many and which of the advertiser's targeting criteria need to match a specific user to determine whether they receive the ad?
    - ii. If it is only a subset, how can I determine which of the criteria used for targeting were matched to me?
  - (d) According to <https://business.twitter.com/en/targeting/tailored-audiences.html>, there are three types of tailored audiences: lists, web, and mobile. I noticed that "Tailored audiences (lists)" and "Tailored audiences (web)" are reflected in the ad-impressions.js file, but are "Tailored audiences (mobile)" also used to target ads? Why does it not appear?
  - (e) When an ad notes that targetingType "Tailored audiences (lists)" was used, does this only mean that I was on the list, or does this also include the "expanding the reach" feature (as explained here: <https://business.twitter.com/en/help/campaign-setup/campaign-targeting/tailored-audiences/TA-from-lists.html>) such that I was not on the list, but only similar to others that were included on the list?
2. The ad-impressions.js file says it contains "Promoted Tweets viewed by the account and associated metadata." This file appears to contain the ads shown in the past 90 days. Does Twitter retain data about ad impressions outside this window?
3. Additionally, I requested the list of "Similar audiences" and "Tailored audiences" from Twitter's settings page, which is sent as a PDF file via email that I have also attached to this request.
  - (a) Is this file accurate for the entirety of a Twitter user's account, 90 days, or some other period of time?
  - (b) Do the "Similar audiences" from this PDF file correspond to the same functionality as "expanding the reach" of a tailored audience, as described at <https://business.twitter.com/en/help/campaign-setup/campaign-targeting/tailored-audiences/TA-from-lists.html>?
  - (c) When are the "Similar audiences" lists created? Does Twitter generate this automatically?

- (d) Does the presence of an advertiser on the PDF’s “Similar audience” list indicate that they explicitly chose to expand the reach of a tailored audience?
- (e) Does an advertiser’s inclusion on our “Similar audiences” list mean that I have seen an ad using that criteria? Or does it simply reflect all advertiser “audiences that are similar to tailored audiences” that I am included in?

Before reporting my concerns to the Information Commissioner’s Office (ICO), I understand that I should give you the chance to respond. You can find guidance on your obligations under information rights legislation on the ICO’s website ([www.ico.org.uk](http://www.ico.org.uk)) as well as information on their regulatory powers and the action they can take.

Please send a full response within one calendar month. If you cannot respond within that timescale, please tell me when you will be able to respond.

If there is anything you would like to discuss, please contact me on the following number [anonymized telephone number].

Yours faithfully,

[Name]

## C.2 Twitter’s Initial Response

Twitter responded with the following on **July 7, 2019**.

Hello,

We found a page in our help center that we think will help you out: (<https://help.twitter.com/en/managing-your-account/accessing-your-twitter-data> and <https://help.twitter.com/en/safety-and-security/privacy-controls-for-tailored-ads>)

If you’ve checked out that page and are still confused, write back to let us know more about where you’re stuck. We’ll do our best to help you out!

Thanks,

Twitter Support

## C.3 Twitter’s DPO’s Response

As described at the beginning of this section, after we received this unsatisfactory response from Twitter Support (“Twitter’s Initial Response”) on **July 7, 2019**, we sent the same request to Twitter’s Data Protection Officer on **July 9, 2019**. Under the GDPR, data controllers are obliged to respond within 30 days from receiving the request. After 31 days, on **August 9, 2019**, Twitter’s DPO responded with the following.

Hello [Name],

Thank you for contacting us.

We are in the process of reviewing your inquiry. Due to its scope, however, we avail ourselves of the deadline extension of 60 days.

If you have any questions about this notice, please let us know.

Sincerely,

Twitter Office of Data Protection

## C.4 Twitter’s DPO’s Detailed Response

On **November 15, 2019**, we received the following, detailed reply.

Hello [Name],

Thank you for your inquiry and patience.

With respect to the questions in your inquiry, we answer in turn below.

1. I have downloaded my data, and in the zip file there is a file called “ad-impressions.js”, which I have also submitted with this request as a courtesy. For each ad in this file, there is a field called “matchedTargetingCriteria”, and I would like to understand how to interpret this information.

- (a) Does matchedTargetingCriteria represent the criteria that the advertiser chose to target this ad? If not, what does it represent?

When you download your Twitter data, an explanatory file is included. This file does indicate that it is the targeting criteria that is used to run the campaign.

- (b) If (a), does the matchedTargetingCriteria shown for a single ad represent ALL the criteria that the advertiser chose for that ad, or are there criteria that the advertiser chose when targeting this ad that are not reflected in the export?

The information represents all of the targeting criteria for how the ad was served to a specific user, in this case @[Twitter handle].

- (c) Do all the matchedTargetingCriteria apply to me, e.g. if an ad is targeted to “Follower look-alikes: @Twitter, Keywords: Privacy, and Locations: United Kingdom”, does that mean that I am seeing this ad because ALL of those criteria apply to me, or only a subset?

Yes, a specific user will be targeted if all criteria match.

(i) If it is only a subset, how many and which of the advertiser’s targeting criteria need to match a specific user to determine whether they receive the ad? (ii) If it is only a subset, how can I determine which of the criteria used for targeting were matched to me?

As mentioned above, they all apply to the user.

- (d) According to <https://business.twitter.com/en/targeting/tailored-audiences.html>, there are three types of tailored audiences: lists, web, and mobile. I noticed that “Tailored audiences (lists)” and “Tailored audiences (web)” are reflected in the ad-impressions.js file, but are “Tailored audiences (mobile)” also used to target ads? Why does it not appear?

The list is a list of device IDs provided by the advertiser. A user may belong to “Tailored audiences (mobile)”, but the response may not be human readable, so these are displayed under the targetingType “Unknown”.

- (e) When an ad notes that targetingType “Tailored audiences (lists)” was used, does this only mean that I was on the list, or does this also include the “expanding the reach” feature (as explained here: <https://business.twitter.com/en/help/campaign-setup/campaign-targeting/tailored-audiences/TA-from-lists.html>) such that I was not on the list, but only similar to others that were included on the list?

It means that the user is on a Tailored audiences list, not similar to others that were included on the list. It does not include the “expanding the reach” feature.

2. The ad-impressions.js file says it contains “Promoted Tweets viewed by the account and associated metadata.” This file appears to contain the ads shown in the past 90 days. Does Twitter retain data about ad impressions outside this window?

In accordance with our Privacy Policy, we do not retain data that is associated with a specific user ID past 90 days. We retain raw logs for 18 months, but they are stripped of the user ID after 90 days.

3. Additionally, I requested the list of “Similar audiences” and “Tailored audiences” from Twitter’s settings page, which is sent as a PDF file via email that I have also attached to this request.

- (a) Is this file accurate for the entirety of a Twitter user’s account, 90 days, or some other period of time?

This is the latest snapshot which includes data no older than 7 days.

- (b) Do the “Similar audiences” from this PDF file correspond to the same functionality as “expanding the reach” of a tailored audience, as described at <https://business.twitter.com/en/help/campaign-setup/campaign-targeting/tailored-audiences/TA-from-lists.html>?

Yes, they are the same.

- (c) When are the “Similar audiences” lists created? Does Twitter generate this automatically?

The lists are automatically created within 24 hours after the audience list is uploaded by the advertiser.

- (d) Does the presence of an advertiser on the PDF’s “Similar audience” list indicate that they explicitly chose to expand the reach of a tailored audience?

Yes, we only expand it when advertisers ask us to do so.

- (e) Does an advertiser’s inclusion on our “Similar audiences” list mean that I have seen an ad using that criteria? Or does it simply reflect all advertiser “audiences that are similar to tailored audiences” that I am included in?

Yes, it reflects the advertiser audiences that are similar to tailored audiences in which you have been included.

Should you have any further questions, please let us know.

Sincerely,

Twitter Office of Data Protection

## D Instructions Provided to Participants For Data Request (Part 1 of the study)

Below is the text we provided to explain to participants how to request their Twitter data. We included detailed and annotated screenshots highlighting each step of the process.

### D.1 Consent Form

Study Title: Twitter Ad Transparency

**DESCRIPTION:** We are researchers at [redacted] doing research to better understand Twitter advertising transparency. In this survey, you will be asked about your experiences and opinions about Twitter. People who are age 18+ and live in the United States or United Kingdom are eligible to participate. Additionally, you must have an active Twitter account. Participation consists of two parts: first, a short 5-minute preliminary survey, and then the main survey, which should take about 35 minutes.

**RISKS and BENEFITS:** The risks to your participation in this online study are those associated with basic computer tasks, including boredom, fatigue, mild stress, or breach of confidentiality. The only benefit to you is the learning experience from participating in a research study. The benefit to society is the contribution to scientific knowledge.

**COMPENSATION:** Participants who complete all tasks will be compensated \$7.86: \$0.86 for Part 1 and \$7.00 for Part 2.

**CONFIDENTIALITY:** No personally-identifiable information will be collected from you. Any reports and presentations about the findings from this study will not include your name or any other information that could identify you. In some cases, you might provide personal stories or beliefs that we might quote or paraphrase as part of our research findings – any personally identifying information will be removed to protect your privacy. We may share the data we collect in this study with other researchers doing future studies – if we share your data, we will not include information that could identify you.

**SUBJECT'S RIGHTS:** Your participation is voluntary. You may stop participating at any time by closing the browser window or the program to withdraw from the study.

[Additional content removed for anonymity]

Please indicate below, that you are at least 18 years old, have read and understand this consent form, and agree to participate in this online research study.

I am at least 18 years old.  Yes  No

I have read and understood this consent form.  Yes  No

I agree to participate in this research study.  Yes  No

### D.2 Introduction and Instructions

Thank you for your participation in our study.

In Part 1 of this study (today), you will log into your Twitter account and request two data downloads. On the next page, you will be guided through the process of requesting your Twitter data.

**NOTE:** The information we collect in this study will not include your personal information. We will NOT ask for your Twitter username, messages, tweets, etc.

We are only interested in data about ads you have seen on Twitter. You will request your entire Twitter data archive today, but in Part 2 of this study, we will provide instructions for uploading only the data we need for our research.

There are two downloads that you need to request in this part of the study.

*How to make the first request:*

- 1) Log into Twitter: <https://twitter.com> (opens in a new tab).
- 2) Click on “More” at the bottom left. On narrower screens, it may only display the icon with three dots, without the word “More”.
- 3) Click “Settings and privacy”.
- 4) Click “Your Twitter data” at the bottom of the menu on the right side.
- 5) If prompted, enter your Twitter password.
- 6) Scroll to the bottom of the page. In the “Download your data” section, click “Request data” in the Twitter row.

If you do not see a button that says “Request data” where the red box appears above, this means you have already requested your data. Continue to the instructions below.

Twitter will email you when your download is ready. There is no need to do anything with this data until Part 2.

To verify that you have successfully requested your data, please copy the text immediately to the left of the “Retrieving data” button, and where the gray box appears in the screenshot below. Paste the text in the text box below.

No text where the gray box appears in the screenshot? You may have previously requested your Twitter data. Instead, please write out the two words on the button that appears instead of the “Retrieving data” button.

*How to make the second request:*

To make the second request, begin with the same first 5 steps.

- 1) Log into Twitter: <https://twitter.com> (opens in a new tab).
- 2) Click on “More” at the bottom left.
- 3) Click “Settings and privacy”.
- 4) Click “Your Twitter data” at the bottom of the menu on the right side.
- 5) If prompted, enter your Twitter password.
- 6) Scroll to the bottom of the page. Now, click “Interests and ads data”.
- 7) Click “Tailored Audiences”.
- 8) Click “Request advertiser list”.
- 9) On the pop-up, click “Request”.

Twitter will email you when your download is ready. There is no need to do anything with this data until Part 2. That’s it for the second request!

To verify that you have successfully requested your data, please copy the text shown where the gray box appears in the screenshot below. Paste the text in the text box below.

Thank you for making the data requests. For today’s last task, please find the summary statistics shown in your Twitter settings on the “Interests and ads data” page.

Please enter the summary statistics into the fields below as numbered in the screenshot. (Fields 1, 2, 3, 4)

### **D.3 Conclusions**

Thank you for completing Part 1 of our study.

It may take a few hours or days until your Twitter data is ready to download.

You will be invited back for Part 2 via Prolific in a few days. In Part 2, you will be given instructions on how to download your Twitter data and upload it to the study.

Part 2 will be a survey that takes 35 minutes to complete.

(Optional) Do you have any final thoughts or comments?

## E Survey Instrument (Part 2 of the study)

This section provides the survey instrument for the main part of our user study.

### E.1 Introduction and General Questions

Thank you for your participation in our study. This survey will take about 35 minutes.

This survey has 4 sections. The first section will ask a few general questions about your data and Twitter.

If a company writes in their privacy policy that “we do not sell your data,” what does that mean to you? In your explanation, please include at least one example of a specific thing you think they would not be allowed to do.

Please rate your agreement with the following statement: I believe that Twitter sells my data.  Strongly agree  Agree  Somewhat agree  Neither agree nor disagree  Somewhat disagree  Disagree

### E.2 Companies

This is the 2nd section (of 4).

In this section, we use the data you uploaded from your own Twitter account. You will be asked about advertisers and up to 4 different advertising methods on Twitter.

The list below shows some companies that showed you an ad on Twitter in the last 3 months.

Please select all of the companies, if any, you remember seeing ads from.

- None of the below
- [Company 1]
- [Company 2]
- [Company 3]
- [Company 4]
- [Company 5]
- [Company 6]
- [Company 7]
- [Company 8]
- [Company 9]
- [Company 10]

### E.3 Section 1 (Targeting Types)

[We repeated this section 4 times for a random selection of 4 *[targeting types]* (e.g., “keywords”) and associated specific *[instances]* of that type (e.g., “my cat is my best friend”) from the participant’s Twitter data. We first asked about the targeting type in the abstract, then about a specific instance from the participant’s Twitter data, and then about the targeting type more generally with a selection of frequent and infrequent instances of that type from the participant’s Twitter data.]

#### E.3.1 Abstract

What does the term **[targeting type]** in the context of online advertising mean to you?

If you have never heard this term before, please write your best guess.

This next section is about *[targeting type]*

*[targeting type]* definition of targeting type

Prior to this survey, I would have expected that advertisers currently target ads on Twitter using *[targeting type]*.  Strongly agree  Agree  Somewhat agree  Neither agree nor disagree  Somewhat disagree  Disagree

#### E.3.2 Specific

On this page, we will give you a specific example of *[targeting type]* from your Twitter data.

According to your Twitter data, you are interested in *[instance]*

Please rate your agreement with the following statements:



I can think of a reason why I Twitter would conclude that I am [interested in, located in or around, would be added to a list of mobile app users by, etc.] [instance].  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

Being [interested in, a speaker of, in the age group, etc.] [instance] describes me accurately.  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

Assume the number of ads you see doesn't change.

I want some of the ads I see to be chosen for me based on being interested in [instance].  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

I am comfortable with Twitter allowing advertisers to target me based on being interested in [instance].  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

### E.3.3 General

Overall, in the last three months, advertisers have targeted up to [#] ads using [targeting type].

In two sentences, please describe your initial reaction to the data above.

This section will ask you to consider how you feel about advertisers using [targeting type] in general. Please rate your agreement with the following statements:

Assume the number of ads you see doesn't change.

I want some of the ads I see on Twitter to be chosen for me using [targeting type].  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

I am comfortable with [targeting type] being used to choose ads for me.  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

I believe it is fair that Twitter allows advertisers to choose ads for me using [targeting type].  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

Please explain your answer to the previous question. If you believe it is fair, why? If you do not believe it is fair, why not?

### E.4 Section 2 (Ad Explanations)

[We repeated this section 6 times for the six ad explanations in randomized order. Each was associated with an ad that the participant had been shown according to their Twitter data, alongside the matched targeting criteria for that ad.]

This is the 3rd section (of 4).

This section will ask for your opinions about potential explanations for why you received a particular ad on Twitter.

To your knowledge, does Twitter have a feature that explains why you received a particular ad?  Yes  No  Don't Know

Imagine a Twitter feature that explains how a particular ad was chosen for you.

In this next section, you will see up to 6 different ads that Twitter has shown you before on its platform, each followed by a different explanation. Then, you will answer questions about what you thought of each ad explanation.

What was the most memorable part of this ad explanation?

What information, if any, did you feel was missing from this ad explanation?

I think this ad explanation shows me all of the information used to target the ad to me.  Yes  No  Don't Know

I feel that this ad explanation was useful.  Yes  No  Don't Know

I feel that this ad explanation gave me enough information to understand how the ad was chosen for me.  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

I would want an ad explanation similar to this one for all ads I see on Twitter.  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

Seeing this ad explanation made me more concerned about my online privacy.  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

Seeing this ad explanation increased my trust in the advertiser who displayed this ad.  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

Do you have any additional comments about this ad explanation?

### E.5 General Opinions About Ad Explanations

If an ad explanation on Twitter did not include all reasons an ad was shown to you, which reasons would be most important for you to see?

For your reference, the previous ad explanations that you've seen in this study will appear below:

Please describe your ideal explanation for ads on Twitter. You are not limited to the things you have seen in this study. Feel free to think big!

## E.6 Demographics

This is the 4th section (of 4). Almost done!

In this section, you will be asked about your Twitter usage and demographics.

Please rate your agreement with the following statement: I believe that Twitter sells my data.  Strongly agree  Agree  Agree  Neither agree nor disagree  Disagree  Strongly disagree

Please explain your answer to the previous question. If you believe Twitter sells your data, why? If you believe Twitter does not sell your data, why not?

What month and year did you join Twitter?

On average, about how many hours do you spend on Twitter each day?  Less than 1 hour  1-2 hours  2-4 hours  4-6 hours  More than 6 hours

Have you ever gone to your Twitter account's settings to look at or make changes to your advertising preferences?  Yes  No  Don't know

Did you look at any of the Twitter files you requested in Part 1 of this study before beginning Part 2?  Yes  No  Don't Know

What is your gender?  Woman  Man  Non-binary  Prefer to self-describe  Prefer not to say

What is your age?  18-24  25-34  35-44  45-54  55-64  65 or older  Prefer not to say

What is the highest degree or level of school you have completed?  Some high school  High school  Some college  Trade, technical, or vocational training  Associate's degree  Bachelor's degree  Master's degree  Professional degree  Doctorate  Prefer not to say

Which of the following best describes your educational background or job field?  I have an education in, or work in, the field of computer science, engineering, or IT.  I do not have an education in, or work in, the field of computer science, engineering, or IT.  Prefer not to say

What is your annual household income?  Less than \$20,000  \$20,000 to \$49,999  \$50,000 to \$99,999  \$100,000 to \$249,999  Over \$250,000  Prefer not to say

When people work on tasks, they are sometimes in situations that can be distracting. How distracted were you while completing this survey?  Not distracted at all  Somewhat distracted  Very distracted

(Optional) Do you have any final thoughts or comments?

## F Twitter’s Explanation About Combining Targeting Criteria

The screenshot shows the Twitter Developer page for 'Targeting Criteria Combinations'. It includes a navigation bar with links for Developer, Use cases, Products, Docs, More, and Labs. The main heading is 'Targeting Criteria Combinations' followed by a sub-heading 'Updated Campaign Workflow'. A paragraph explains that campaigns can target broadly with geo, gender, language, and device/platform criteria, and that additional criteria like interests and keywords can be added. A key point is that if no targeting criteria is specified, the line item will target all users worldwide. A table lists 'Primary' and 'Other' targeting types. Below the table, it explains how criteria are combined using logical operators (U, AND, OR) and provides a geo example and additional examples.

| “Primary” Types    | Other Types           |
|--------------------|-----------------------|
| Followers          | Locations             |
| Tailored Audiences | Gender                |
| Interests          | Languages             |
| Keywords           | Devices and platforms |
| TV                 | Age                   |

Targeting criteria will be combined for your ad group such that:

- “Primary” Targeting Types will get **U**’d (i.e. put in a logical union).
- Other Targeting Types will get **AND**’d.
- Same types will get **OR**’d.

### Some examples

At a glance: **[(Followers) U (Tailored Audiences) U (Interests) U (Keywords)] AND (Location) AND (Gender) AND (Languages) AND (Devices and Platforms)**

A Geo example:

Let’s say we want an ad group for our campaign to serve targeting:

- Twitter users in the U.S., England, and Canada (Location)
- who are Women (Gender)
- derived from Tailored Audiences list (“Primary”)
- with Keywords (“Primary”)

The targeting criteria will be:

**[US OR GB OR CA] AND [Female] AND [Tailored Audiences U Keyword]**

### Additional examples

- Select Gender and Geo but no primary: **(Male) AND (US OR GB)**
- Select Gender, Geo, Interest: **(Female) AND (CA) AND (Computers OR Technology OR Startups)**
- Select Gender, Geo, Interest, Tailored Audiences, Keywords: **(Male) AND (GB) AND (Cars U Tailored Audiences for CRM U autocross)**

Figure 6: Twitter’s explanation on their developers page about how targeting criteria are combined.

**What's the policy?**

Advertisers using keyword targeting in timeline may not select keywords that target sensitive categories. Unless otherwise provided in the country-specific requirements below, this policy applies globally.

**Which sensitive categories may not be targeted?**

- Alleged or actual commission of a crime
- Health
- Genetic and/or biometric data
- Negative financial status or condition
- Political affiliation or beliefs
- Racial or ethnic origin
- Religious or philosophical affiliation or beliefs
- Sex life
- Trade union membership

**How does this policy vary from country to country?****U.S.**

Advertisers targeting the U.S. may target based on trade union membership.

**What do advertisers need to know about this policy?**

As with all advertising platforms, there are certain obligations to follow when using Twitter for advertising. Review our guidelines and make sure you understand the requirements for your brand, business, promoted content, and targeting criteria. You are responsible for all your promoted content and targeting on Twitter. This includes complying with applicable laws and regulations regarding online advertisements.

When configuring the keywords in your campaign, be aware of the audience(s) you may reach and align your targeting and message in a way that is appropriate and compliant with our [Twitter Ads Policies](#). Using keywords to target users based on sensitive categories is not permitted in certain countries, could be considered inappropriate or offensive, and may reflect poorly on your brand, product, or service. Targeting users based on sensitive categories is a violation of our Twitter Ads policies.

Twitter takes violations of its [Twitter Ads Policies](#), the [Twitter Rules](#), and [Terms of Service](#) seriously. We will examine reported violations and take appropriate action, which may include removal of offending advertisements and advertisers from the Twitter Ads platform.

Figure 7: Twitter's policy for prohibiting targeting based on sensitive categories from <https://business.twitter.com/en/help/ads-policies/other-policy-requirements/policies-for-keyword-targeting.html>.

## G Targeting Types Seen by Participants

Table 4: The number of participants who saw each targeting type in the survey.

| Targeting Type | # Participants | Targeting Type | # Participants |
|----------------|----------------|----------------|----------------|
| Location       | 79             | Tailored web   | 58             |
| Age            | 77             | Tailored lists | 56             |
| Lookalikes     | 67             | Conversation   | 51             |
| Language       | 66             | Mobile         | 45             |
| Platform       | 64             | Event          | 40             |
| Gender         | 63             | Movie/TV       | 40             |
| Keyword        | 63             | Retargeting    | 37             |
| Interest       | 59             | Behavior       | 33             |

## H Regression Tables

In this section, we present tables depicting the full results of our mixed-effects ordinal logistic regression models analyzing targeting type data from Part 1 of the user study.

Table 5: Mixed-effect ordinal logistic regression model of how participants’ agreement responding to **General: Fair** (“I believe it is fair that Twitter allows advertisers to choose ads for me using *targeting type*”) varied by targeting type.

| Factor                      | Baseline | Odds Ratio | $\beta$ | Std. Error | $z$    | $p$   |
|-----------------------------|----------|------------|---------|------------|--------|-------|
| <b>Type: Age</b>            | Interest | 0.640      | -0.446  | 0.389      | -1.146 | .252  |
| <b>Type: Behavior</b>       | Interest | 0.155      | -1.863  | 0.475      | -3.924 | <.001 |
| <b>Type: Conversation</b>   | Interest | 0.362      | -1.017  | 0.431      | -2.360 | .018  |
| <b>Type: Event</b>          | Interest | 0.401      | -0.914  | 0.464      | -1.969 | .049  |
| <b>Type: Gender</b>         | Interest | 0.350      | -1.050  | 0.408      | -2.576 | .010  |
| <b>Type: Keyword</b>        | Interest | 0.310      | -1.171  | 0.404      | -2.898 | .004  |
| <b>Type: Language</b>       | Interest | 4.480      | 1.450   | 0.429      | 3.494  | <.001 |
| <b>Type: Location</b>       | Interest | 0.262      | -1.339  | 0.388      | -3.449 | <.001 |
| <b>Type: Lookalikes</b>     | Interest | 0.218      | -1.522  | 0.396      | -3.845 | <.001 |
| <b>Type: Mobile</b>         | Interest | 0.128      | -2.057  | 0.441      | -4.668 | <.001 |
| <b>Type: Movie/TV</b>       | Interest | 0.386      | -0.952  | 0.460      | -2.068 | .039  |
| <b>Type: Platform</b>       | Interest | 0.497      | -0.699  | 0.410      | -1.706 | .088  |
| <b>Type: Retargeting</b>    | Interest | 0.409      | -0.895  | 0.458      | -1.954 | .051  |
| <b>Type: Tailored lists</b> | Interest | 0.061      | -2.802  | 0.430      | -6.511 | <.001 |
| <b>Type: Tailored web</b>   | Interest | 0.120      | -2.118  | 0.418      | -5.064 | <.001 |

Table 6: Mixed-effect ordinal logistic regression model of how participants’ agreement responding to **General: Comfortable** (“I am comfortable with *targeting type* being used to choose ads for me”) varied by targeting type.

| Factor                      | Baseline | Odds Ratio | $\beta$ | Std. Error | $z$    | $p$   |
|-----------------------------|----------|------------|---------|------------|--------|-------|
| <b>Type: Age</b>            | Interest | 0.743      | -0.297  | 0.394      | -0.755 | .451  |
| <b>Type: Behavior</b>       | Interest | 0.162      | -1.817  | 0.489      | -3.714 | <.001 |
| <b>Type: Conversation</b>   | Interest | 0.326      | -1.121  | 0.438      | -2.556 | .011  |
| <b>Type: Event</b>          | Interest | 0.439      | -0.823  | 0.451      | -1.823 | .068  |
| <b>Type: Gender</b>         | Interest | 0.525      | -0.644  | 0.411      | -1.569 | .117  |
| <b>Type: Keyword</b>        | Interest | 0.290      | -1.239  | 0.404      | -3.071 | .002  |
| <b>Type: Language</b>       | Interest | 5.411      | 1.689   | 0.440      | 3.838  | <.001 |
| <b>Type: Location</b>       | Interest | 0.305      | -1.188  | 0.391      | -3.039 | .002  |
| <b>Type: Lookalikes</b>     | Interest | 0.251      | -1.383  | 0.396      | -3.495 | <.001 |
| <b>Type: Mobile</b>         | Interest | 0.110      | -2.204  | 0.450      | -4.897 | <.001 |
| <b>Type: Movie/TV</b>       | Interest | 0.433      | -0.838  | 0.457      | -1.836 | .066  |
| <b>Type: Platform</b>       | Interest | 0.399      | -0.918  | 0.408      | -2.250 | .024  |
| <b>Type: Retargeting</b>    | Interest | 0.286      | -1.253  | 0.459      | -2.730 | .006  |
| <b>Type: Tailored lists</b> | Interest | 0.063      | -2.758  | 0.434      | -6.363 | <.001 |
| <b>Type: Tailored web</b>   | Interest | 0.158      | -1.846  | 0.422      | -4.375 | <.001 |

Table 7: Mixed-effect ordinal logistic regression model of how participants’ agreement responding to **General: Want** (“Assume the number of ads you see doesn’t change. I want some of the ads I see on Twitter to be chosen for me using *targeting type*”) varied by targeting type.

| Factor               | Baseline | Odds Ratio | $\beta$ | Std. Error | z      | p     |
|----------------------|----------|------------|---------|------------|--------|-------|
| Type: Age            | Interest | 0.304      | -1.190  | 0.377      | -3.157 | .002  |
| Type: Behavior       | Interest | 0.152      | -1.885  | 0.477      | -3.955 | <.001 |
| Type: Conversation   | Interest | 0.390      | -0.941  | 0.421      | -2.235 | .025  |
| Type: Event          | Interest | 0.403      | -0.910  | 0.443      | -2.055 | .040  |
| Type: Gender         | Interest | 0.301      | -1.202  | 0.398      | -3.024 | .002  |
| Type: Keyword        | Interest | 0.221      | -1.511  | 0.392      | -3.852 | <.001 |
| Type: Language       | Interest | 3.318      | 1.200   | 0.415      | 2.893  | .004  |
| Type: Location       | Interest | 0.254      | -1.369  | 0.377      | -3.633 | <.001 |
| Type: Lookalikes     | Interest | 0.200      | -1.611  | 0.382      | -4.217 | <.001 |
| Type: Mobile         | Interest | 0.099      | -2.312  | 0.435      | -5.311 | <.001 |
| Type: Movie/TV       | Interest | 0.406      | -0.900  | 0.444      | -2.028 | .043  |
| Type: Platform       | Interest | 0.210      | -1.561  | 0.393      | -3.975 | <.001 |
| Type: Retargeting    | Interest | 0.273      | -1.299  | 0.445      | -2.917 | .004  |
| Type: Tailored lists | Interest | 0.061      | -2.803  | 0.420      | -6.680 | <.001 |
| Type: Tailored web   | Interest | 0.114      | -2.171  | 0.406      | -5.345 | <.001 |

Table 8: Mixed-effect ordinal logistic regression model of how participants’ agreement responding to **Specific: Comfortable** (“I am comfortable with *specific example of targeting type* being used to choose ads for me”) varied by targeting type.

| Factor               | Baseline | Odds Ratio | $\beta$ | Std. Error | z      | p     |
|----------------------|----------|------------|---------|------------|--------|-------|
| Type: Age            | Interest | 0.745      | -0.295  | 0.354      | -0.833 | .405  |
| Type: Behavior       | Interest | 0.266      | -1.324  | 0.447      | -2.963 | .003  |
| Type: Conversation   | Interest | 0.276      | -1.287  | 0.401      | -3.212 | .001  |
| Type: Event          | Interest | 0.271      | -1.307  | 0.434      | -3.011 | .003  |
| Type: Gender         | Interest | 0.686      | -0.377  | 0.381      | -0.990 | .322  |
| Type: Keyword        | Interest | 0.306      | -1.183  | 0.372      | -3.176 | .001  |
| Type: Language       | Interest | 6.017      | 1.795   | 0.401      | 4.477  | <.001 |
| Type: Location       | Interest | 0.394      | -0.932  | 0.359      | -2.598 | .009  |
| Type: Lookalikes     | Interest | 0.356      | -1.033  | 0.359      | -2.877 | .004  |
| Type: Mobile         | Interest | 0.081      | -2.511  | 0.425      | -5.909 | <.001 |
| Type: Movie/TV       | Interest | 0.284      | -1.258  | 0.433      | -2.903 | .004  |
| Type: Platform       | Interest | 0.575      | -0.553  | 0.373      | -1.482 | .138  |
| Type: Retargeting    | Interest | 0.188      | -1.673  | 0.434      | -3.852 | <.001 |
| Type: Tailored lists | Interest | 0.121      | -2.115  | 0.406      | -5.210 | <.001 |
| Type: Tailored web   | Interest | 0.105      | -2.255  | 0.399      | -5.655 | <.001 |

Table 9: Mixed-effect ordinal logistic regression model of how participants’ agreement responding to **Specific: Want** (“Assume the number of ads you see doesn’t change. I want some of the ads I see on Twitter to be chosen for me using *specific example of targeting type*”) varied by targeting type.

| Factor               | Baseline | Odds Ratio | $\beta$ | Std. Error | z      | p     |
|----------------------|----------|------------|---------|------------|--------|-------|
| Type: Age            | Interest | 1.315      | 0.274   | 0.340      | 0.806  | .420  |
| Type: Behavior       | Interest | 0.746      | -0.293  | 0.429      | -0.683 | .495  |
| Type: Conversation   | Interest | 0.348      | -1.054  | 0.387      | -2.725 | .006  |
| Type: Event          | Interest | 0.237      | -1.439  | 0.428      | -3.362 | <.001 |
| Type: Gender         | Interest | 1.016      | 0.016   | 0.359      | 0.044  | .965  |
| Type: Keyword        | Interest | 0.401      | -0.913  | 0.364      | -2.510 | .012  |
| Type: Language       | Interest | 8.434      | 2.132   | 0.380      | 5.619  | <.001 |
| Type: Location       | Interest | 0.788      | -0.239  | 0.347      | -0.688 | .491  |
| Type: Lookalikes     | Interest | 0.436      | -0.831  | 0.352      | -2.358 | .018  |
| Type: Mobile         | Interest | 0.064      | -2.750  | 0.433      | -6.357 | <.001 |
| Type: Movie/TV       | Interest | 0.331      | -1.106  | 0.418      | -2.644 | .008  |
| Type: Platform       | Interest | 0.965      | -0.036  | 0.355      | -0.102 | .919  |
| Type: Retargeting    | Interest | 0.175      | -1.742  | 0.435      | -4.002 | <.001 |
| Type: Tailored lists | Interest | 0.237      | -1.440  | 0.379      | -3.797 | <.001 |
| Type: Tailored web   | Interest | 0.138      | -1.984  | 0.389      | -5.100 | <.001 |

Table 10: Mixed-effect ordinal logistic regression model of how participants’ agreement responding to **Specific: Accurate** (“*Specific example of targeting type describes me accurately*”) varied by targeting type.

| Factor               | Baseline | Odds Ratio | $\beta$ | Std. Error | $z$    | $p$             |
|----------------------|----------|------------|---------|------------|--------|-----------------|
| Type: Age            | Interest | 2.939      | 1.078   | 0.339      | 3.182  | <b>.001</b>     |
| Type: Behavior       | Interest | 0.477      | -0.740  | 0.384      | -1.926 | .054            |
| Type: Conversation   | Interest | 0.437      | -0.827  | 0.356      | -2.326 | <b>.020</b>     |
| Type: Event          | Interest | 0.348      | -1.056  | 0.392      | -2.696 | <b>.007</b>     |
| Type: Gender         | Interest | 5.158      | 1.641   | 0.380      | 4.315  | <b>&lt;.001</b> |
| Type: Keyword        | Interest | 0.428      | -0.848  | 0.333      | -2.548 | <b>.011</b>     |
| Type: Language       | Interest | 9.691      | 2.271   | 0.392      | 5.789  | <b>&lt;.001</b> |
| Type: Location       | Interest | 1.880      | 0.631   | 0.334      | 1.889  | .059            |
| Type: Lookalikes     | Interest | 0.492      | -0.710  | 0.324      | -2.192 | <b>.028</b>     |
| Type: Mobile         | Interest | 0.104      | -2.266  | 0.399      | -5.684 | <b>&lt;.001</b> |
| Type: Movie/TV       | Interest | 0.501      | -0.692  | 0.401      | -1.724 | .085            |
| Type: Platform       | Interest | 2.922      | 1.072   | 0.342      | 3.138  | <b>.002</b>     |
| Type: Retargeting    | Interest | 0.165      | -1.801  | 0.406      | -4.434 | <b>&lt;.001</b> |
| Type: Tailored lists | Interest | 0.240      | -1.428  | 0.347      | -4.118 | <b>&lt;.001</b> |
| Type: Tailored web   | Interest | 0.146      | -1.923  | 0.364      | -5.282 | <b>&lt;.001</b> |

Table 11: Mixed-effect ordinal logistic regression model of how participants’ agreement responding to **Specific: Reason** (“I can think of a reason why [*phrase explaining that Twitter would conclude that I am similar to, or described by, specific example of targeting type*]”) varied by targeting type.

| Factor               | Baseline | Odds Ratio | $\beta$ | Std. Error | $z$    | $p$             |
|----------------------|----------|------------|---------|------------|--------|-----------------|
| Type: Age            | Interest | 1.696      | 0.528   | 0.329      | 1.604  | .109            |
| Type: Behavior       | Interest | 0.434      | -0.835  | 0.396      | -2.108 | <b>.035</b>     |
| Type: Conversation   | Interest | 0.406      | -0.902  | 0.363      | -2.481 | <b>.013</b>     |
| Type: Event          | Interest | 0.327      | -1.117  | 0.403      | -2.772 | <b>.006</b>     |
| Type: Gender         | Interest | 2.253      | 0.812   | 0.362      | 2.243  | <b>.025</b>     |
| Type: Keyword        | Interest | 0.458      | -0.782  | 0.354      | -2.210 | <b>.027</b>     |
| Type: Language       | Interest | 8.172      | 2.101   | 0.396      | 5.308  | <b>&lt;.001</b> |
| Type: Location       | Interest | 1.742      | 0.555   | 0.333      | 1.665  | .096            |
| Type: Lookalikes     | Interest | 0.523      | -0.649  | 0.332      | -1.956 | .050            |
| Type: Mobile         | Interest | 0.203      | -1.595  | 0.395      | -4.040 | <b>&lt;.001</b> |
| Type: Movie/TV       | Interest | 0.582      | -0.542  | 0.401      | -1.352 | .176            |
| Type: Platform       | Interest | 3.055      | 1.117   | 0.352      | 3.169  | <b>.002</b>     |
| Type: Retargeting    | Interest | 0.306      | -1.184  | 0.410      | -2.892 | <b>.004</b>     |
| Type: Tailored lists | Interest | 0.287      | -1.248  | 0.360      | -3.462 | <b>&lt;.001</b> |
| Type: Tailored web   | Interest | 0.248      | -1.393  | 0.367      | -3.795 | <b>&lt;.001</b> |