Audience strategy

SEGMATIC

Expert Series #5

The comprehensive guide to capitilising on the biggest trend in paid search campaigns

Contents

1. Audiences are on the rise

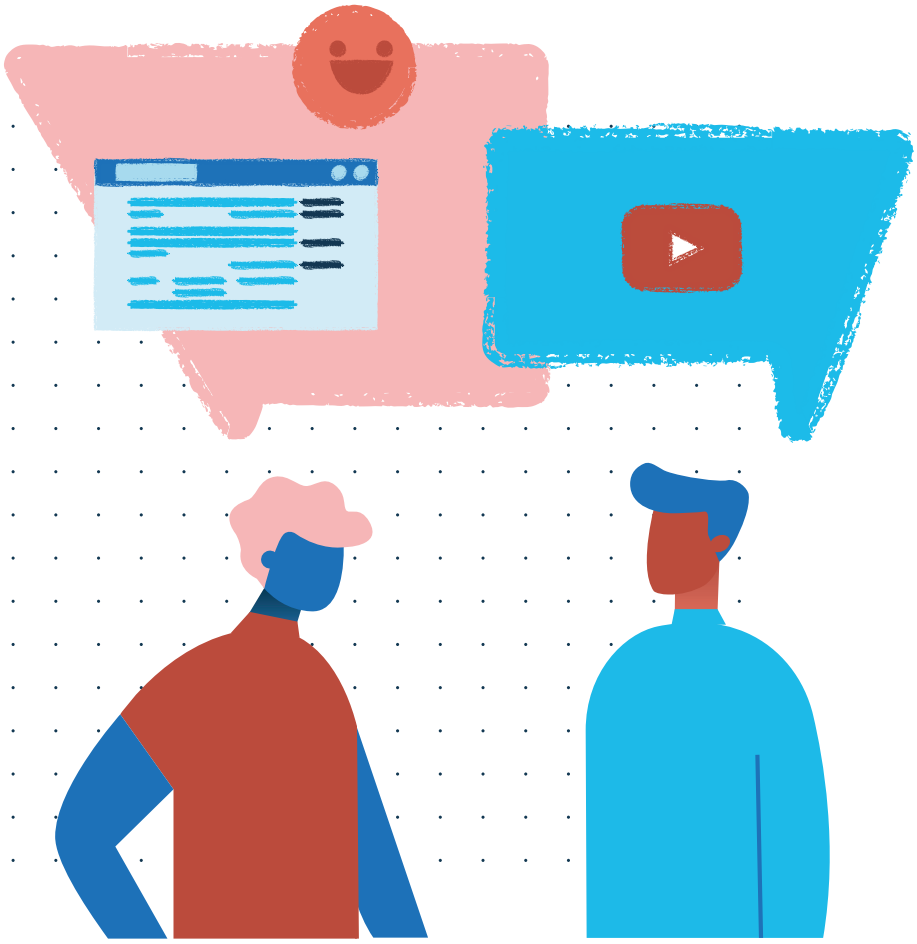

The search term is the greatest targeting mechanic ever created, and the power of the search term has driven Google's success. But, audiences are challenging the dominance of the search term.

For advertisers, keywords have long been the best way to tap into the power of search terms and reach the right customer with the right ad. That's changing fast, and the rise of audiences is one of the biggest trends in paid search advertising.

This shift is happening because Google has been busy collecting data about its users. What used to be limited to country, language, and device is now a vast and growing body of information that advertisers can layer into their targeting.

Google has also been rolling out more and better functionality to allow advertisers to use this data. For a long time, Facebook led the way on this, but Google is aggressively playing catch-up.

As a result, audiences are exploding. Audience strategy is now at least as important as keyword strategy.

How you can make the most of audiences

At Segmatic, we've responded to this shift by developing tools and approaches that take full advantage of new and emerging audience technology.

Up until recently, advertisers could put almost all their energy into keyword strategy; paying a little bit of attention to the person behind the search; and using the small amount of data that Google gave us on things like location and device type.

Simply put, our approach is to take advantage of the many types of user data that are available to us now; and to do this while maintaining single-keyword ad groups, and keeping control of ad copy and bidding for individual keywords.

In this guide, we're going to explain a bit about how audiences have changed; then we'll unpack our approach and explain in detail how we execute it.

2. Types of audience data

There are lots of ways to track users and use their behaviour to optimise how we advertise to them.

We'll start by looking at how this data is collected, then move on to how we segment this data and use it for targeting. The data is collected in a couple of ways: the pixel, Customer Match, and Google data.

The pixel

The classic one is the pixel. This is a small piece of code, usually JavaScript, that sends information to Google about how users are interacting with a website. The client puts the pixel on their website, and creates rules about what data to collect and how to organise it.

Typically, we'll take this first party data from a client and segment it in a couple of ways. We might create a data set with information on all of the people who visit their site. Or, we might segment it more finely: broken down to sets like people who viewed the shirts page but not the suits page, or people who got as far as the checkout but didn't make a purchase.

This is all useful information, but it can be tricky to work with because there tends to be overlap between the groups.

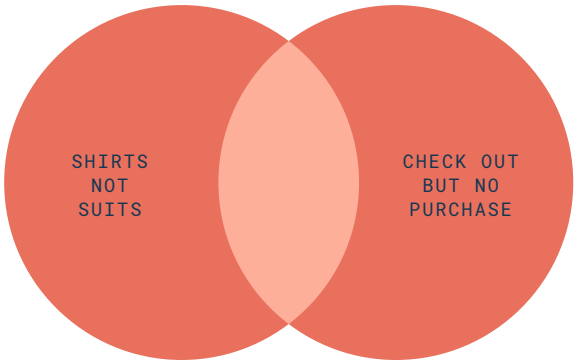

Customer Match

Customer Match is a simple way to collect and use audience data. We collect data on people who sign up and give us their email address or phone number. This gives us a certain amount of data on the user, which we can segment in lots of different ways. Once it's segmented, we upload it into Google Ads Editor and use it to inform our targeting.

Since the introduction of GDPR, there are more limits on what we can do with Customer Match marketing, but it's still an effective channel.

Google data

This is the big one, and it's getting bigger as Google expands the suite of audience data it offers advertisers. This data comes in three main forms, each of which we'll explore here: demographic data, similar audiences, and affinity audiences.

This is the big one, & it's getting bigger.

Demographic data

Google offers advertisers user profiles that they construct by combining their own user data with data that they buy from external sources. Using regression analysis and educated guesswork, they use this data build profiles.

The result is data on things like age and gender, as well as new categories like household income.

This type of data can be incredibly helpful for targeting specific groups of users. And it's easy to work with because there's no overlap between groups in any given category. Everyone fits into one age bracket and one income bracket (or the 'unknown' bracket if Google hasn't figured you out).

Income brackets

UNDER 25,000 PER ANNUM	OVER 25,000 PER ANNUM	UNKNOWN

That clean definition between brackets is extremely powerful from an analytics perspective; and it's sadly missing from similar or affinity audiences.

Similar audiences

If you're familiar with Facebook's Lookalike Audiences, you'll understand how Google's similar audiences works. The two products are very similar, and the Google offering is pretty great.

Similar audiences gives you a big group of Google users who share characteristics with a smaller set of your own users that you've curated. You use your own data to create a set of people; you give that to Google; and they create a big set of people who resemble the people in your set.

It's helpful to understand a little about how Google makes this projection from the specific to the general. It's based on search history data. Google looks at the pixel and identifies when a user typically takes a particular action. For example, two hours before they buy a photo book, people searched for `photo printing`, `bananas`, and `photo books`.

Because Google have huge volumes of data to work with, they can they identify common features between users and groups of users, including—and this is borderline creepy—features that predict things those users are likely to do in the future.

Affinity audiences

If similar audiences are the custom version, affinity audiences are off-the-shelf, ready-made audience sets that Google curates themselves. An example is people who, based on their search history, may be about to buy a car. Google has created data sets like these because they know that lots of people are interested in them.

If you sell something more obscure than cars, Google probably hasn't created an affinity audience for it (yet), so similar audiences are your best option.

3. Applying audiences to ad groups

Once you have these segmented audiences, there are two different ways you can apply them to ad groups: observation (previously called bid only) or targeting (previously called target and bid).

Observation means that everyone who searches for your keyword will see the same ad, whether or not they're in your audience. However, you can see data for the subset of people who searched for the keyword and are also in your audience, and you can enter a different bid for these users.

Targeting means that in order for your ad to show to a user, that person must search for your keyword and be in a particular audience. If they searched for the keyword but they're not in your audience, your ad won't show.

Overlap is a dangerous thing

As we mentioned earlier, the beauty of Google's demographic data is that it's contiguous—there's no overlap between groups. Each individual person fits neatly into one of the age brackets, one of the income brackets, and so on, with no spill-over.

Similar and affinity audiences are messier. Within these data sets, people fit into multiple categories. You might have an audience that's made up of people who searched for suits, people who searched for shirts, and an overlap group of people who searched for both.

Our solution is to duplicate the campaign and negative the overlapping audiences from each other.

This might not sound that serious, but it means you lose control.

First off, you lose control over which users see your ads. If you were to apply an audience that contains overlaps to a neat single-keyword ad group, you get situations like Google picking one audience over another in an unpredictable way that's based on your bidding but beyond your control.

Our solution when this happens is to duplicate the campaign and negative the overlapping audiences from each other, wherever this is possible.

4. The bidding problem

The real problem isn't splitting impressions: it's bidding. It is extremely difficult to optimise bidding for any section of an audience that overlaps with another.

This is especially true if the overlap group is different from the two original groups.

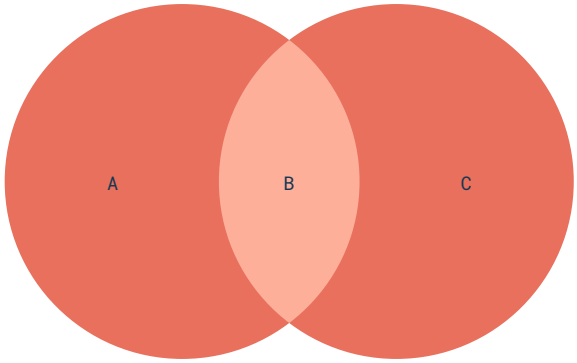

If the overlap group C has characteristics that are significantly different to those of A or B, it's really hard to optimise your bidding for C.

The bidding death spiral

We'll start by looking at a realistic example of this problem and the wrong way to tackle it, and then we'll explain how we've solved it.

The problem

Imagine you have two audiences: Audience 1 and Audience 2. You want to enter a different bid for each audience, which we'll call High Bid and Low Bid. There's an overlap between the two audiences, which we'll call Audience X

First attempt

For the sake of our example, let's say that:

- The people who only fall in Audience 1 have a conversion rate of 5%
- The people who only fall in Audience 2 have a conversion rate of 5.1%
- The people in Audience X, the overlap, have a conversion rate of 1%

A lot of advertisers would allow the overlap to be split down the middle, tossing each half into one of the other audiences; so that one half of Audience X goes into Audience 1, making Basket 1, and the other half goes into Audience 2, making Basket 2. To begin with, we follow this approach and we enter the same bid for both cohorts.

This approach is simple, but it foolishly ignores the fact that Audience X has an average conversion rate that's different from those of Audience 1 and Audience 2. In reality, there is almost always a variance in conversion rate between these types of audience.

The result of this first try is that Basket 2 performs better, thanks to the slightly higher conversion rate of Audience 2.

Second attempt

Based on this result, you tweak your approach. Basket 2 did better, so you bid it up.

This shifts the distribution: all of the people in Audience X move over into Basket 2. This pushes the conversion rate for Audience 2 below that of Audience 1.

Third go

You might guess what happens next: our advertiser reacts to this development by bidding down Audience 2 and bidding up Audience 1. And they keep doing this switching until they retire. The continuous switching makes reporting and tracking a nightmare, so the rest of their career is played out in agony.

How to end the cycle

Alternatively, our advertiser might realise after a few rounds that they don't really know what's going on with Audience X, and they can't really analyse this group or optimise their bidding for it.

This is bad because it means they can never bid in a way that differentiates between Audiences 1 and 2; so they end up bidding the same on both, which is not ideal.

This is bad because you can never bid in a way that differentiates between audiences.

How big a problem this is depends on the scale of the overlap. If the overlap is less than 1%, it probably doesn't matter. If it's bigger, you need to find a way to isolate Audience X as best you can, and bid on them separately.

So that's what we did.

5. Our solution

There are really two problems here: the first is that there's overlap between our data sets; and the second is that the overlap is messing with our bidding strategy. We solve both problems by eliminating, as far as possible, any overlap between audiences.

For first party data, we use Google's tools to split the data into neat isolated subsets with no overlap. This takes a bit of work, but if you're disciplined it can be done. For example, we can create one audience made up of everyone who visited the site minus all the people who viewed shirts, and a second audience that's just people who viewed shirts.

For similar and affinity audiences, some overlap is unavoidable because we can't see the regression analysis that Google did when they put the data sets together. However, there are things we can do to minimise the effects of this overlap.

The benefits of going big

First, we minimise this overlap by focusing on big, homogenous audiences. An audience made up of people who searched for `shirts` is less likely to create overlaps than a lot of smaller audiences who searched for `white shirts`, `blue shirts`, etc.

We only split a big audience up if we decide that there's a big advantage to showing different ad copy to different sections. This means more work and it increases the chance of overlap; so we only do it where there's a big upside to showing different ads.

We might, for example, split out new customers and existing customers, or men and women when there's evidence that they behave very differently in response to a certain keyword.

The decision tree

Then we apply the decision tree: this is a technique we use when we're building accounts to decide when to create a new single-keyword ad-group. Like most trees, it's made up of two questions—do we want to bid differently, and do we want to show a different ad?

The bidding question doesn't apply to audiences: because audiences are layered on top of a keyword, you can choose to bid up or down within an audience without impacting your underlying keyword structure.

So, the question we ask is whether we want to show different ad copy to one of your audiences. If the answer is yes, then we duplicate the campaigns and negative them from each other. If Campaign 1 and Campaign 2 are on the same keyword, we make Audience 1 positive in Campaign 1 and negative in Campaign 2, and vice versa. This gives us different ads for different audiences.

This creates some extra work—two campaigns to manage instead of one—so we only do it if there is a big differential between the performance of the two ads.

The most common example of a big differential is between new and existing customers. You might want to offer much more aggressive promotions to acquire new customers.

By segmenting our audiences in this way, we can take full advantage of the audiences data Google now offers, while keeping control over our ads and bidding.