Account structure

SEGMATIC

Expert Series #1

The comprehensive guide to building bigger, smarter accounts that dominate the auction

Contents

1. Start with great account structure

Great account structure provides the foundation for everything we do at Segmatic—smarter bidding, better ads, control at scale, everything.

And we expect our accounts to achieve great things. In this guide, we'll explain our thinking on account structure, and give you a step-by-step guide to how we build, maintain, and optimise accounts.

Great ads rely on great account structure; and we want every ad to appear to the customer as if it was custom-written by their own personal ad copywriter. An ad should respond not just to their search term but to the searcher's intent, reflecting whatever was in their head when they entered the search term.

To this end, the best ad copy mirrors the search term. Get that right, and the searcher will feel that they've found what they're looking for; and people who feel that way are far more likely to click through and buy something.

To achieve this, our accounts are made up of tens of thousands of unique ads, and our ad structure makes sure that the right ad is triggered for the right person every time.

On top of matching intent, we also want to dominate the results page. When we get our account structure right, this happens:

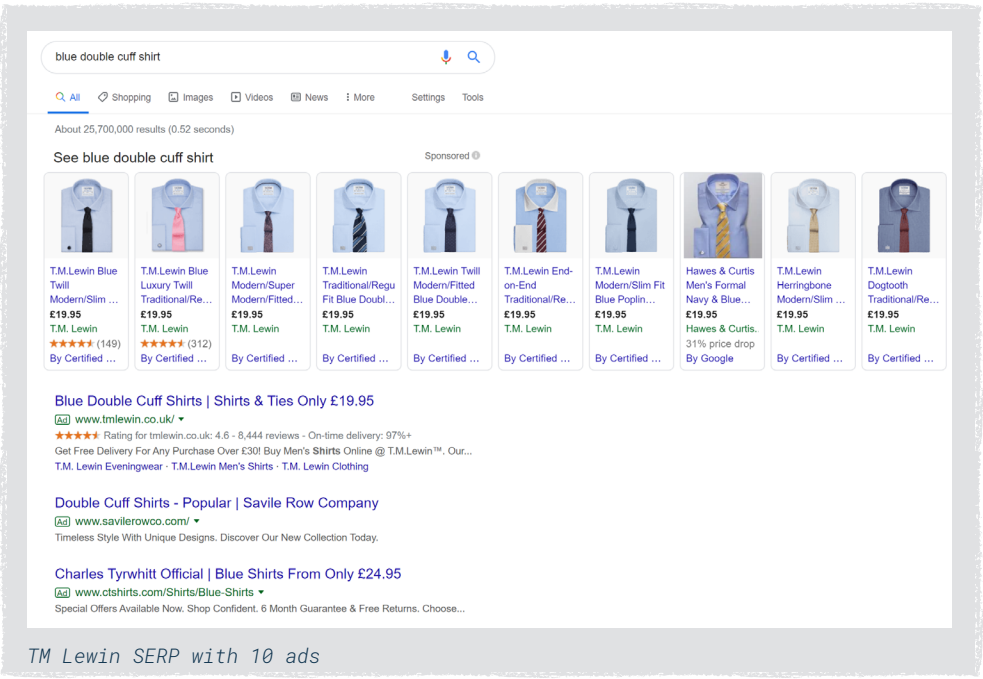

TM Lewin SERP with 10 ads

These results are built on a structure that offers absolute, granular control over every aspect of our accounts, down to keyword level.

This lets us understand how each ad is performing, quickly and cleanly tweak individual ads and ad groups, and set up intelligent feedback loops that mean that our accounts learn and evolve over time, essentially optimising themselves. And we want to do all of that at enormous scale, with no loss of control as we scale up to hundreds of thousands of ad groups.

To do this, we deisgned an approach to account structure that makes all of this possible. Our building blocks are lists of products and product attributes, which we cross-match to create thousands of individual, isolated ad groups. If you were to look under the hood of one of our accounts it would look pretty complicated, and it is. But, once you understand the step involved, it all becomes clear.

2. Our account structure philosophy

Our starting point for account structure is the simple fact that keywords are not the same as search terms. People sometimes use the two interchangeably, but most of the time they're not the same thing.

A search term is the word or collection of words that people use to search—literally what they type into the little box on Google. Keywords are the individual words that make up a search term, so they are usually a subset of search terms. Google ads are triggered by keywords, not by search terms. People type search terms into Google, but advertisers bid on keywords.

At Segmatic, we structure accounts to achieve as close a relationship as possible between search term and keyword, no matter what the customer has searched for. We're constantly amazed by the way users phrase their search queries, and we constantly adapt our tactics to respond to the search terms they use.

Building on this, we have three guiding principles:

- **Create isolated units of search terms**
  Our accounts are made up of unique, isolated units of search terms; each of which we can observe and optimise in isolation.

- **Use the decision tree to decide when to create isolated units**
  The decision tree is a series of binary decisions that tell us how granular we should be in creating our isolated units.

- **Use the right match types to bridge between search terms and keywords**
  Using exact match and broad match modified (BMM), we can narrow the gap between the keywords we're bidding on and the search terms people are typing into Google.

How does this philosophy work in practice?

We apply this thinking to a wide range of accounts for different clients in different industries, and we find that it consistently delivers a huge number of very small wins that add up to really extraordinary gains. It works because it lets us write better, more targeted ads, and lets us boost click-through rate, conversion rate, and ultimately profits.

Next, we want to unpack our approach in detail, moving step-by-step from philosophy to implementation, beginning with our obsession with control.

3. Control is the Holy Grail

The quality we value above all others is control. We want to control who sees the ad, what ad they see, and when they see it. We also want to understand at a granular level how each of our ads are performing and why, and to use this understanding to optimise performance.

Choose the right match type

The first key to control is to use the right match type. Google offers advertisers four match types. From broadest to narrowest, they are broad, broad match modified, phrase, and exact. The broader match types allow let more in: so they allow for more potential traffic, but it's less targeted traffic. The narrower match types are more focused: so they won't trigger on as many search terms. But when they do, we can be confident that our keywords will be more relevant to the searcher.

We get the greatest level of control by combining exact and broad match modified (BMM). BMM casts a wider net than exact, and it makes sure that our ads appear in response to search terms that wouldn't trigger on exact match. Over time, as we get to understand an account better, we shift more and more of our campaigns to exact match.

Create isolated units

Isolation is at the heart of our account structure. We can only truly achieve control when every ad group is an isolated, single-keyword unit. This means that we can observe, adjust, and optimise each ad group in isolation.

We create isolated units by layering campaigns on top of each other and then using negative matching to keep them separate.

We can only achieve control when every ad group is an isolated, single-keyword unit.

Without isolated units, things get messy: a search term could trigger multiple keywords and multiple ads, or a single search term can appear in multiple campaigns. When this happens, it's hard for us to know which campaigns are performing well and which are not; and it means you won't have a reliable data set as a basis to optimise the performance of our campaigns.

This sort of messiness makes it pretty much impossible to scale your campaigns without complicating and messing up your account structure. With isolated units, scaling is simple: you can add whole new type of product groups or sets of modifiers while maintaining the shape and performance of your account structure.

Our approach takes time and effort, and to do it right you need to understand your product, your customers, their search behaviour, and the wider environment your business is operating in. The good news is that creating isolated units also create a virtuous circle. They produce usable insights about your campaigns and your business as a whole; and some of these can help you to make better decisions in areas well beyond paid search advertising.

4. Oh wise decision tree, should we make a new isolated unit?

You're probably wondering how big these units are. A frequent and important question for us is when should we keep keywords grouped together in the same isolated unit, and when do we split them off into a unit of their own? To answer that question, we consult the decision tree.

Like most trees you'll come across in your life, the decision tree consists of two questions:

- Can I write a better ad for this keyword?
- Should I bid differently on this keyword?

If the answer to either one is yes, then we split it out into its own isolated unit.

Can I write a better ad?

Let's say you're running advertising on the keyword +SHIRTS, and you see in your search term analysis the term `14.5" shirts`. You need to decide whether to split `14.5" shirts` out into its own unit.

So, can you can write a better, more effective ad for `14.5" shirts`?

In this case, the answer is probably yes. An ad for 14.5" shirts is going to be more effective for customers searching for `14.5" shirts` than a generic one. So, we split.

Should I bid differently?

You need to decide whether to split `men's shirts` and `men's shirts online`, or leave them in the same unit. Your ad copy will be the same for both, so they fail the 'Can I write a better ad?' test.

However, you should probably be bidding differently on them. We know that someone who adds the word online to a search term is signalling stronger intent to purchase than someone who searches for the product name alone. We also know that this intent leads to a different enough conversion rate that we will bid differently on it. The answer is yes, you should bid differently: so, we split it out into its own ad group.

That, in a nutshell, is the decision tree.

5. How we build an account

Now for the fun part, let's build ourselves an account!

Here we'll take you step-by-step through how we build an account. For the purposes of this exercise, we're selling three products: shirts, suits, and ties.

An online advertiser walks into a bar, pub, tavern, hostelries...

With any campaign, our first step is to create lists of keywords that we want to bid on. These are the building blocks. We create detailed lists because we want to show one ad to a person who searches for `mens shirts`, and a very different ad to someone who searches for `shirts online`, and a different ad again to someone who searches for `mens shirts online`.

For this we use two types of keyword list: Product Lists, and Modifier Lists.

A Product List (PL) is pretty much what it sounds like: it's a list of keywords that name things we want to sell. For this example, we have three items; so we create three Product Lists:

PRODUCT LISTS	
PL 1	Shirts
PL 2	Suits
PL 3	Ties

A Modifier List (ML) is a list of the non-product keywords that appear in search terms, like blue or striped. These keywords often give us more detail on the type of product the searcher is interested in; and, as an added bonus, they can tell us something about the searcher themselves. In our example, and in general, these Modifier Lists fall into four categories.

MODIFIER LISTS	
ML 1	Gender: men's, women's...
ML 2	Intent: buy, purchase, online, web...
ML 3	Discount: sale, code, voucher...
ML 4	Characteristics: blue, white, twill, 14", 14.5"...

Gender tells us what type of shirt the searcher is looking for—men's or women's, girl's or boy's. We can use this information to trigger the right ad, and stuff the ad copy with the gender-specific keyword.

Intent signals to us how likely it is that the searcher plans to make a purchase: someone searching for `buy shirts` has different intent to someone searching for `shirts`. The first person is likely to buy a shirt; the second might be doing research for their novel on the shirt business.

Discount signals the potential value that the searcher represents: words like discount or sale signal a lower value customer.

Characteristics describe something about the product that's being searched for. In our suits, shirts, and ties example, characteristics will be things like blue and double-breasted. Just like gender, these help us to trigger an ad that's tailored to the searcher's specific interest.

Where do Product Lists & Modifier Lists come from?

It's simple enough to create Product Lists. Once we know what we're selling, and what words people are using to describe that product when they search, we can put together a list of synonyms. Over time, drawing on search term analysis, we build out this list.

For Modifier Lists, we start by making a list based on information from the client and our own research. On day one of a campaign, we might have a Modifier List for shirts that includes all of the types of fabric that our client uses to make shirts. And just like Product Lists, we can add to our Modifier Lists through regular search term analysis.

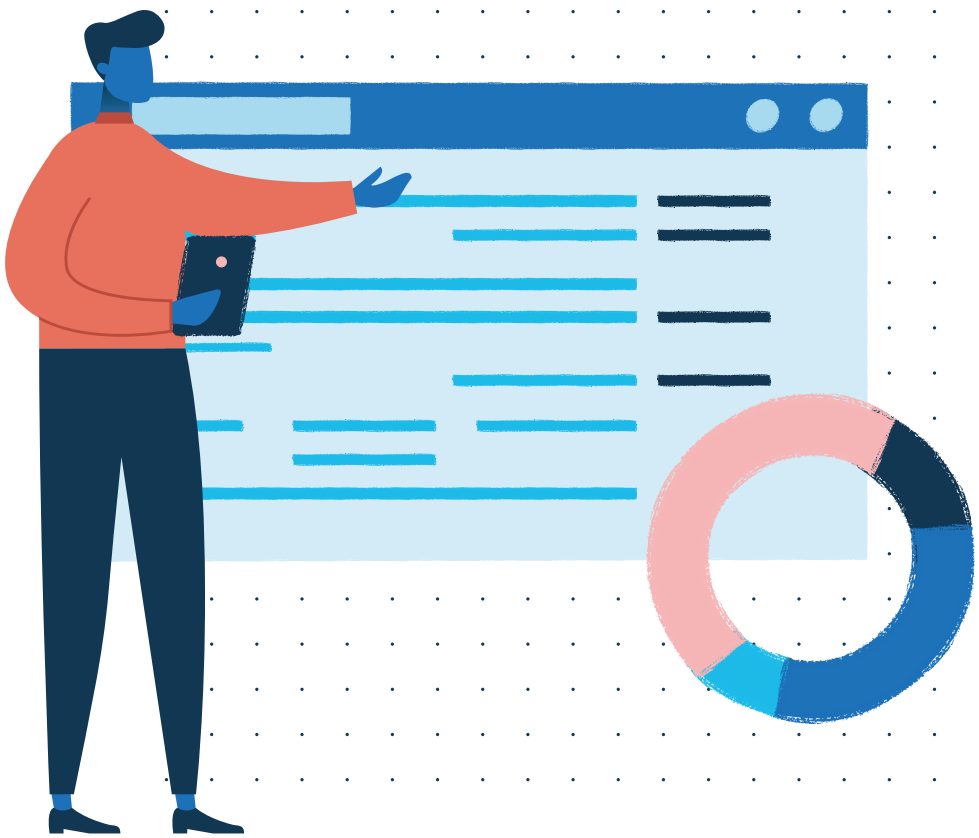

Proactive clue-based search term analysis

At Segmatic, we have a unique approach to search term analysis. From day one, we analyse the actual search terms that people are using, and mine them for keywords to add to our lists. We might find, for example, that a lot of people are searching for `polyester shirts` are being served an ad for our client and clicking on it, only to find that our client doesn't sell polyester shirts.

We make a decision on what to do with keywords we discover through our search term analysis, based on the keyword's conversion rate. If the keyword is leading to conversions, we add it to a product or Modifier List. If it's not leading to conversions, we add it to a negative match list so we're not wasting budget bidding on it.

Most advertisers take this sort of a reactive approach to their search term analysis. We take it a big step further, employing a proactive clue-based approach.

First, we scour the data for search terms that suggest broader trends; then we use these clues to make projections and take action today on things that we can predict are going to appear in our search term data in the future.

As an example, we have a client that sells business shirts online. Looking at their search term data, we spotted the term `Manchester United shirt`. That's a type of shirt our client does not sell. The approach most people would take is to this is to negative out those keywords and leave it at that. In a few months, they might see `Arsenal shirt` in the data and negative that out as well. We call this the reactive approach.

Our approach is to see `Manchester United shirt` as a clue: it tells us that people are searching for, and will continue to search for, that type of shirt; and they'll continue to click on our ad, which loses us money. So once we identified the search term was a clue, we made a list of all of the football teams in the world, and most of the other professional sports teams, and we negatived the whole list out of our campaigns.

For more on this, see our guide to search term analysis (Expert Series #2).

When a Product List & a Modifier List love each other very much...

Product Lists and Modifier Lists are our building blocks; but, the real magic happens when we combine them to create Product Modifier Lists, or PMLs. Here's how this works.

PMLs layer campaigns on top of each other to create isolated units, like so:

PML 1	PL 1		
PML 1	PL 1	ML 1	
PML 1	PL 1	ML 2	
PML 1	PL 1	ML 4	
PML 1	PL 1	ML 4	ML 4
PML 1	PL 1	ML 1	ML 2

Each Product Modifier List is made up of many individual campaigns.

Let's zoom in on PML 4, which is layering Product List 1 (Shirts) with Modifier List 4 (Characteristics).

Let's say for the sake simplicity that for shirts, ML 4 is made up of 8 characteristics:

- white
- blue
- striped
- check
- cotton
- linen
- single cuff
- double cuff

Based on this, we will build out 8 individual ad groups, with different bidding and ad copy for each one.

In its most basic form, this is how we build an account and how we ensure that keywords adhere as closely as possible to search terms. This ensures our ads appear custom made.

It gets more complex when you mix in the three other types of Modifier List (Gender, Intent, and Discount), when you mix exact match and broad match modified, and when you introduce negative matching.

The real beauty of this account structure is that it gives you nice, clean data that you can use to optimise each individual campaign; and that it allows you to learn by testing hypotheses and experimenting on specific campaigns.

6. Fun with negative matching

To achieve isolated units, we use negative matching in a couple of ways that you may not have seen before: we call these forcing the match and the negative hierarchy.

Forcing the match

This is employed at times when you're bidding on both exact and BMM, but you want the keyword (when left unmodified by the searcher) to match on exact match only. You might do this because you know that your conversion rate is higher on exact than it is on BMM. We force the match by applying the exact matches as negatives for our BMM campaigns.

Don't let Google decide which of your ads is triggered

Say someone searches for `shirts`; presumably that person needs a shirt, and they're keeping their search term short and sweet because they haven't got a minute to spare.

You're bidding on the keyword SHIRTS in both exact match and BMM. In other words, you're live on both [SHIRTS] and +SHIRTS.

In this scenario, Google will decide at random which campaign is triggered by which search term: some will trigger on exact match and some on BMM, and you can't control when this happens. We really don't like giving up this control to Google—both on principle, and because we should never trust that Google's incentives one hundred percent overlap with our client's. We take control of the choice by using a negative match to make sure that every time the exact term is searched for, only the exact match campaign is triggered.

There's a right way and a wrong way to force the match

Sadly, you can't flip a switch in Google Ads to make sure it triggers on exact match rather than BMM: so what we do is we negative match all exact matches from BMM. By doing this, you force Google to trigger only on exact match.

The old-fashioned and slightly lazy way of doing this, which we still see quite often, is to set a higher bid on exact match across the board so it triggers first on exact and never triggers on BMM. In our view, this is bad practice because it messes up your bidding. We go to a lot of effort to make sure that we're bidding precisely the right amount in every auction we enter, and we're not willing to monkey with that so we can trigger on the right match type. If our bidding formula tells us we should be bidding the exact same amount on exact match and BMM, then that's what we bid. Forcing the match lets us solve this problem without artificially distorting our bids, and that makes us happy.

The negative hierarchy

Say you've created an account made up of PMLs and it's produced tens of thousands of unique ads. For any given search term, there will be a couple of front-runners for the best ad to trigger; and you want the most appropriate, effective ad to appear every time. For us, the perfect ad means ad copy that mirrors the search term as closely as possible.

We make sure this happens by creating a hierarchy of ads and using negative matching to trigger the best one. This is the negative hierarchy.

Let's look at an example: someone searches for `buy men's shirts`. Our client sells shirts, so we're live on some or all of the keywords in that search term. We enter the auction, and our ad appears.

But which ad? We have two options: an ad with the headline 'Buy shirts' or one with 'Men's shirts'. We know that 'Buy shirts' is better—more likely to lead to a sale; so we want to make sure that the 'Buy shirts' ad is triggered.

We should note that this is a very simple example with an A or B choice with an obvious winner. In reality; it's rarely as straightforward as this.

Our approach to this is to decide on a hierarchy—a determination of which type of ad is best, second best, third best, and so on, in a given situation. We then use negative matching to make sure that the right ad is triggered in response to the right search term.

The nuts and bolts of how this works

Let's go back to the example we used when we were building an account. We started out with Product Lists and Modifier Lists:

PRODUCT LISTS	
PL 1	Shirts
PL 2	Suits
PL 3	Ties

MODIFIER LISTS	
ML 1	Gender: men's, women's...
ML 2	Intent: buy, purchase, online, web...
ML 3	Discount: sale, code, voucher...
ML 4	Characteristics: blue, white, twill, 14", 14.5"...

And we brought them together in many different combinations to create product Modifier Lists.

PRODUCT MODIFIER LISTS	PRODUCT LISTS	MODIFIER LIST 1	MODIFIER LIST 2
PML 1	PL 1		
PML 2	PL 1	ML 1	
PML 3	PL 1	ML 2	
PML 4	PL 1	ML 4	ML 4
PML 5	PL 1	ML 4	ML 2
PML 6	PL 1	ML 1	

You can see in this matrix that PML 1 only includes PL 1, which we know from the PL is 'shirts'.

PML 5 is more complicated: it's made up of PL1 (shirts), ML4 (characteristics), and ML2 (intent). In turn, each of the MLs is made up of lots of individual keywords; so ML4 is a big list of all of the relevant colours, sizes, patterns, fabrics, and many other qualities of shirts.

Characteristics are king

Going back to our 'Buy men's shirts' example, we know that if someone searches for `buy mens' shirts` we want to trigger a 'Buy shirts' ad rather than a 'Men's shirts' ad. So, in this case, we say that purchasing intent (ML2) is more important than gender (ML1). It's higher up the hierarchy.

Similarly, let's say someone searches for `buy 14.5" shirts`, do we want to trigger an ad with the headline 'Buy shirts' or one with '14.5" shirts'? This time the answer is '14.5" shirts'. So, we conclude that characteristics (ML4) is higher on the hierarchy than purchasing intent (ML2).

Through asking a few more questions and drawing some conclusions from the answers, we end up with a hierarchy for each Product List that tells us which ML is most important, second most important, and so on. It might look like this:

RANK	MODIFIER LIST
1	ML4—Characteristics
2	ML2—Purchasing Intent
3	ML1—Gender
4	ML3—Discount

How we apply this hierarchy to our accounts

We use negative matching to make sure that ads that correspond to the lists at the top of the hierarchy are triggered, and those further down are not.

As with forcing any match, there isn't a switch you can flick to tell Google Ads to trigger ML4 ads over ML2 ads, so you have to force this using layered negative matching.

ML1	ML2	ML3	ML4
ML2	ML4	ML1	
ML4		ML2	
		ML4	

We've identified that we want ML1 to trigger ahead of ML2 and ML4; so we negative match ML2 and ML4 in any Product Modifier Lists that include ML2. In our example above, that's PML3 and PML5.

For any PMLs that don't include MLs, like PML1 in our example, we negative out all MLs if they're active.

Bulk negativing

There are a few things that we negative match out of our accounts across the board. Brand terms are a big one: our clients' names and the names of their competitors all get negative matched. If someone types in the name of your brand, they should go into a custom-made brand campaign, no matter what else is in the search term. Similarly, if someone types in the name of a competitor, they go into a competitor campaign.

What's the pay-off?

We apply the negative hierarchy across all of our accounts. If you do this right, it brings a couple of benefits. It boosts click-through rate, and therefore profits. It also lets you isolate the ad, which gives more control. Better click-through rate leads to an improved quality score—an added bonus that makes us between 10% and 11% happier.

Negative matching to avoid overlap

We use negative matching to keep our accounts tidy—making sure that our isolated units are truly isolated. We also use negative matching to suppress our own ads on certain keywords. Pretty much everyone uses negative matching to make sure they're not bidding on keywords that have nothing to do with their business.

At Segmatic, we also use negative matching on keywords that we are live on as a way of keeping things tidy within our campaigns. Let's say we have two keywords: +SHIRTS and +MEN'S +SHIRTS. When someone searches for `men's shirts`, we have two options: trigger the broad, general ad we've written for +SHIRTS, or the more specific, targeted one for +MEN'S SHIRTS?

To make sure +MEN'S +SHIRTS triggers every time, we negative match +MENS out of the +SHIRTS ad group.

This negative matching has a couple of other advantages. Without it, the search term `men's shirts` would trigger for both of them. This is bad news for a couple of reasons. First off, we never want to let Google decide which ad appears. They'll choose one at random or based on their own preferences—not good. Negativing removes the risk that two people searching for the same search term could see different ads. And, it ensures that we're not splitting data across different campaigns and ad groups, which would make it a lot harder to analyse. Finally, the negatives mean that +MEN'S +SHIRTS now has more volume, giving us more data on men's shirts for us to analyse.

Now you know all our account structure secrets

At its core, this is all about control over our accounts at a granular level: keeping our campaigns in clean, isolated units and generating high quality, usable data that we can use to create feedback loops that optimise campaigns performance. We believe in this approach because we've seen the results it produces for our clients (and so have they!). We hope you see the value too and can learn from our work to start building bigger, smarter accounts that dominate the auction.