



An Introduction to A/B Testing



Table of Contents

An Introduction to A/B Testing	3
Understanding A/B Testing	4
What is A/B Testing?	
How does A/B Testing work?	
A/B Testing Terminology	
The Core Principles of A/B Testing	
A/B Testing vs. Multivariate Testing	
Getting started with A/B Testing	7
Designing A/B Tests	15
Building Test Variations	
Best Practices for A/B Test Design	
Avoiding Common Testing Mistakes	
Running A/B Tests	20
Implementing tests correctly	
Splitting Traffic	
Test Duration	
Monitoring Tests and Collecting Data	
Maintaining Statistical Significance	
Analyzing and Interpreting Results	22
Interpreting Data Accurately	
Deriving Actionable Insights	
Using Artificial Intelligence To Drive A/B Tests	23
Generating Hypotheses	
Automating Audience Allocation	
Determining Statistical Significance	
Recommending Changes to Future A/B Tests	
Creating Visualizations	



An Introduction to A/B Testing

In marketing, optimization is critical, but often difficult to achieve. With A/B testing, you can improve the effectiveness of marketing variables (such as copy, images, layouts, and so on) by creating different versions and seeing which one converts better.

You might create two different images on an ad, send out two versions of a newsletter, or test out multiple website layouts. By conducting this experiment, you can not only improve conversion rates, but better understand your audience to create better content.

In other words, if you're not A/B testing, you're missing out on:

- ✔ **Gaining an advantage over your competitors**
- ✔ **Learning critical insights about your audience**
- ✔ **Improving your bottom line**

Understanding A/B Testing

What is A/B Testing?

A/B testing, also known as split testing, is a way to test two types of content against each other to see which performs better. Typically, an A/B test changes one variable at a time—such as the copy or image—with the goal of picking the one that generates the most conversions. In fact, it's one of the most popular ways to optimize conversion rates, with 77% of businesses **leveraging** it to identify areas of improvement.

Though many use A/B testing to optimize web pages, you can use it to improve marketing variables such as:

- ★ Ads
- ★ Newsletters
- ★ Blogs
- ★ Freebies
- ★ Apps
- ★ Websites

How does A/B Testing work?

Let's say you want to run an ad for your business. You create the copy and image for it, but before running it, you want to A/B test it. To do so, you change only the copy and run two versions of the ad: A and B.

After it runs for sometime, you discover B significantly outperforms A. You decide to shut off the A variant and invest in the high-performing ad.

Though simple, this hypothetical is A/B testing in a nutshell. You can even extend this testing further—you have the optimal caption, what about the image, the audience demographics, and so on? You can A/B test any variable, so long as you do it one at a time.

A/B Testing Terminology

Throughout this guide, we use different terms to describe A/B testing. If you plan on running A/B tests, it's worth knowing the terminology so you can explain it to others in your business.

- ✔ **Control:** the variable that remains constant throughout the experiment (say, the color of a CTA button when A/B testing the text over it)
- ✔ **Hypothesis:** an informed guess you make based on previous research and information
- ✔ **Experiment:** the procedure of testing your hypothesis (let's say you hypothesize changing the copy a certain way will improve conversion rates)
- ✔ **Conversion rate:** the number of visitors to your content that convert into customers



The Core Principles of A/B Testing

In this section, we cover the must haves for an effective A/B test. By incorporating these principles, you more easily measure results and determine which variation performs better.

01 Conduct one test on one variable at a time

As the example demonstrates, you should pick one variable to test at a time. Otherwise, you muddy the results. Let's say, for example, you decide to run the A/B test by changing the copy *and* the image. If A performs better than B, what contributes to that: the copy or the image? You can't know, unless you test one variable at a time.

02 Make tiny changes, too

Even small changes can impact your conversion rates. So, while bigger changes will significantly alter performance, remember to make tinier ones, too. For example, changing the color of a CTA button from green to red could significantly impact conversion rates.

03 Consider larger variables

Your variable can be as big or as small as you like, from the color of a CTA button to the entire marketing element. You could design, for example, two completely different landing pages or emails and see which one performs best. This yields the biggest changes, so it may make sense to start here before tweaking the smaller elements.

04

Look beyond conversion rates

A/B testing directly affects conversion rates, but in the long run, it can have a significant impact on your bottom line. Whenever possible, try tracking how A/B testing impacts your leads, demo requests, sales, and so on so you can really see the effects of A/B testing.

05

Use two audiences of the same size and background

Like all experiments, you want to control as many variables as possible. In this case, make sure the audiences you conduct the A/B test on are similar in size and background. Let's say you want to A/B test a newsletter. To ensure the audience stays the same, randomize them into two separate groups.

05


Test at the same time

Seasonality can impact the performance of your marketing efforts. Remember to also control the variable of timing to ensure the accuracy of your A/B tests. If you run test A in one month and test B in the next, you might not know whether that difference played a major role in performance.

A/B Testing vs. Multivariate Testing

Both A/B and multivariate testing aim to optimize performance. However, multivariate testing, as the term suggests, involves testing many variables simultaneously. This means conducting a series of simultaneous A/B tests rather than one at a time, often for high-traffic sources.

Multivariate testing is typically performed by more technical marketers, as it's a fairly complex process to set up. Though this guide focuses primarily on A/B testing, much of the information can be applied to multivariate testing should you pursue it.



Getting started with A/B Testing

If you've not experimented before, don't fret — in this section, we lay out a step-by-step guide on starting your first A/B test. Remember to keep in mind the principles laid out in the previous section.

01 Decide what you want to test

You can A/B test a million different things. But to get started, pick one marketing variable you want to focus on. Do you want to improve your online advertising efforts? Or do you want to increase conversion rates on a landing page?

02 Define your objective

Ask yourself — *what do I hope to gain by A/B testing?* A higher conversion rate, of course, but more than that, what insights do you want to glean from your test? An A/B test on a landing page, for example, could provide insights into how your audience interacts with your website, whereas an ad might show how they engage with advertising.

03 Formulate your hypothesis

Your hypothesis acts as the bridge between the experiment and the results. Think of it as a conjecture you make based on experience, previous research, and even your gut. Your hypothesis might look like this in the end:

If we add a banner image to our landing page, **then** the conversion rate will increase.

04 Identify your testable elements

As you build landing pages, calls-to-action, and email campaigns, you've probably wondered about the elements you can test and optimize to increase conversion rates. Should you change the background color of the landing page? How can you modify the language on the call-to-action to attract the most clicks? What if you removed all images from your email campaign?

Not all variables are created equal, and some may prove more worthy of your time than others. Below is a list of areas to focus on for landing pages, emails, and calls-to-action.

Elements That You Can Optimize on a Landing Page

★ Offers

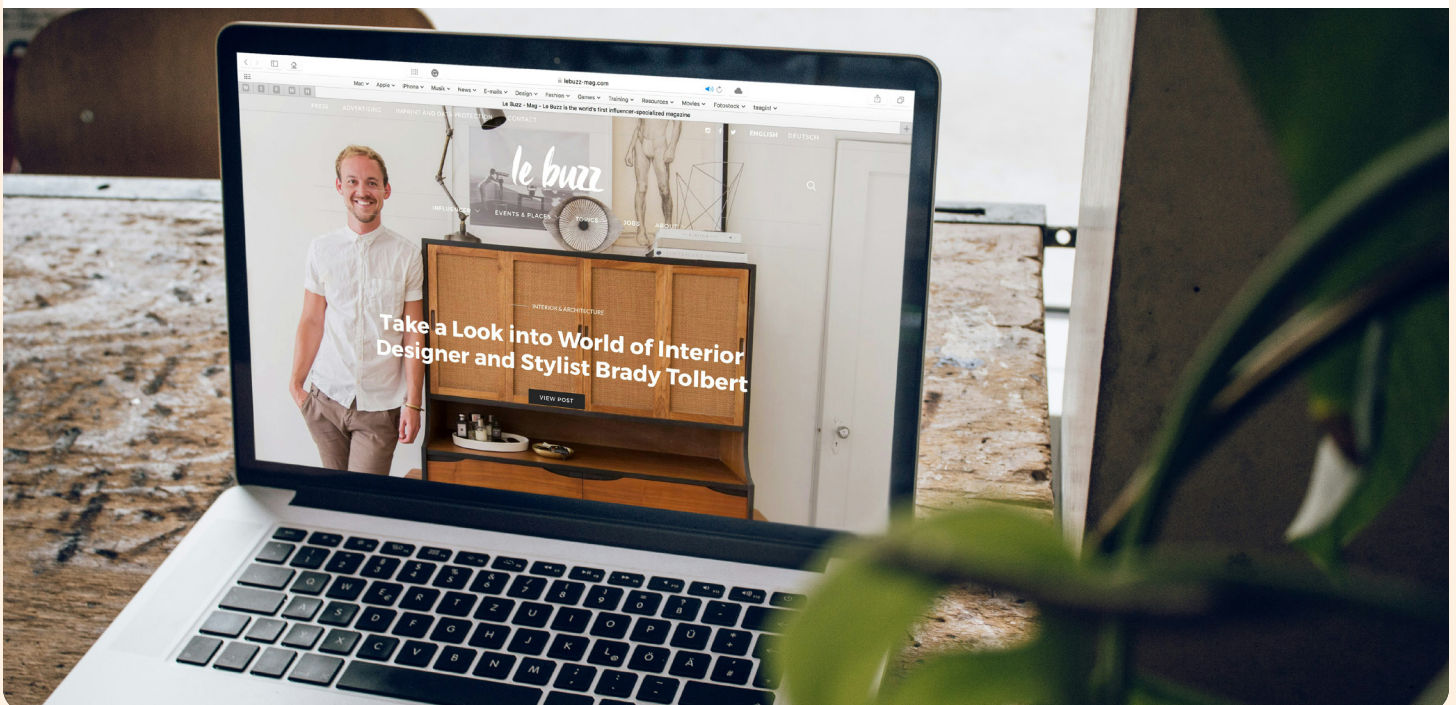
You should start your optimization process by finding out what types of offers convert the most visitors into leads and which offers help you push leads down the sales funnel.

Examples of offers to test include ebooks, webinars, discount codes, coupons, sales, demos, and more. For instance, at HubSpot, we found that ebooks perform better than webinars in converting visitors to leads, but webinars do better than ebooks in converting leads to customers. That has led us to maintain a balanced mix of content types.

More middle-of-the-funnel offers, such as sales consultations and demos, will most likely be characterized by high customer close rate. You can also test different topics and find out how they compare in driving business results.

★ Copy

How should you position your offer? What messaging will entice your reader? Should you add testimonials to strengthen the visitor's incentive? People looking for information online will pay attention to the description of your offer. Consider different landing page copy that can help you drive more conversions. For instance, bullet points and data-driven content have traditionally performed well for us at HubSpot. Start with a radical test in which you compare a short, one-paragraph long description to long, but still valuable copy.





Form Fields

Should your lead capture form only request an email address or should it ask for more information?

Inbound marketers diverge in their decision whether to place content behind a form or keep it form-free. Some argue that forms create friction in the lead generation process, while others believe forms are essential for qualifying traffic and prioritizing work for the sales organization.

Landing page A/B testing enables you to evaluate how your audience reacts to different questions -- what prospects are willing to answer and what information they would rather not share. Form fields help you qualify leads and nurture them. With form-free content, you have to relinquish more lead nurturing control — you didn't ask for their information, so how can you get back in touch with them unless they bookmark your site?


You can also test the placement of your form fields. For example, using an [exit-intent lead flow](#) on your pages or fully gating your content.



Whole Page

As we mentioned earlier, taking the entire page as the page variable is the fastest way to achieve drastic results and produce a landing page that drives many conversions. It's also a great approach when you're not seeing gains from micro optimizations (like changing the color, image, or copy).

Make iterations to the whole page that affect image placement, form length, and copy. Once you have a statistically significant result pointing to the variation that performed better, you can continue optimizing through smaller tweaks.



Elements That You Can Optimize on a **Call-To-Action**



Placement

The argument over where the “best” place to put a CTA is never ending. Some say the best place to put a CTA is in the top left hand of a page. After all, we start reading from left to right and from top to bottom.

However, at HubSpot, we’ve found that different assets (including blog posts, emails, landing pages, and more) have seen different results for which CTA placement performs best. That’s why we’re constantly testing our placement over time. Try A/B testing the right and left side of the page. Later, try testing in-text CTAs vs. traditional CTAs. You can even try pop-up, exit flows, and more.



Size

The size of a CTA is tightly related to the context of the page and the other characteristics of your call-to-action. For instance, the CTA will naturally be large if it includes a graphic or an image that strengthens the message. Create an A/B test to see if a big call-to-action that adds value to the message—e.g. a customized blog CTA—attracts more clicks than its control.



Remember: bigger CTA’s will help draw attention to the desired action you want your user to take, but if it’s too big, it can also overpower the content and decrease conversion rates.



Color

What’s the ideal color for CTAs? Should you use a bold, in-your-face color like bright red or should you focus on brand consistency and determine the color based on the design of the page the CTA is on?

There’s no one clear answer. This can only be solved for each individual brand through A/B testing. The goal of a CTA is to stand out and draw your users’ eyes so they take your desired action. Make sure to use a contrasting color from the rest of your page, but run A/B tests to determine which colors do a better job of capturing users’ attention. Pay attention to your brand recognition, however. Choose colors that make sense for your brand and website, not just colors that you think will be the loudest.



★ Copy

The copy of your CTA should be short, sweet, and to the point. But it should also effectively describe what the user will get if they take your intended action. A/B testing copy is an effective way to test what kind of copy resonates best with your audience. Maybe it's copy that uses social proof to convey the impact of an offer. Or, perhaps it's bulleted copy that describes the details of what's inside the offer. Only one way to find out: A/B test it.

★ Graphic

While you should focus on your call-to-action text, don't forget that graphics can help convey meaning and strengthen your message. Experiment with various shapes besides the standard button-like CTA form and test how the new look affects your click-through rate.

Elements That You Can Optimize In Your Email

Many of the elements to test and optimize we have discussed so far overlap across different channels. Offers, copy and image are certainly some of the variables you want to always keep in mind. They apply to email marketing as well, so we won't repeat them here.

★ Format

What's the best way to lay out your email content to drive the most engagement? Should your message be structured as a newsletter, digest, or dedicated send?

A/B testing can help determine the right format of your email marketing campaigns. For instance, newsletters can perform well in spreading the news about a few different pieces of information, such as events, new offers, discounts and event product announcements.

Dedicated sends, on the other hand, can help you drive the most conversions to one call-to-action. Run an A/B test to determine which format yields the results you seek to achieve.

★ Layout

Similar to testing formats, layout is another element that you can test and optimize for on an ongoing basis. Experiment with different image and call-to-action placements. Look at not only your clickthrough rates, but also your conversions to determine which layout is most effective.

★ Timing

Naturally, the optimal timing for sending marketing emails will vary by industry and even company. Should your prospect be sent your next email one hour later or one day later? Identify the best time to convert prospects into leads by running split test based on your audience.



★ Sender

Back in 2011, we at HubSpot ran A/B email marketing tests that showed messages coming from a personal sender name receive higher CTR rate than messages from a generic sender name.

So, instead of sending emails from Marketing Team, we started to make the senders of our marketing emails real experts from our marketing team. To reinforce this personalization and create consistency, we even added the person's signature at the end of the message and included an image, their title, and social links.

Personalization is a well-known best practice in email marketing. But you should still test sender names with your audience. Ensure that a change like this yields positive results for your company before you implement it.

★ Subject Lines

Subject lines are the part of emails that can grab your recipients' attention immediately and convince them to open the email and read more. You really want to get the subject line right.


Try testing the tone of the message, using personalization, asking questions, and more.

If you decide on testing email subject lines, try to deduct a lesson from your experiment. Is it the length of the subject line, mention of a discount, or certain formatting (e.g., brackets or colon) that made the difference?

★ Target Group

Who are you sending that email to and why will that particular audience find the content valuable? Segmentation can help you get high response rates because your recipients will find the messages particularly valuable.

That will also increase the chances of them sharing your email with friends and coworkers.



05

Determine your sample size and duration

In many ways, an A/B test pushes against your day-to-day operations. So, when running one, first determine your sample size and duration. Do you want to leverage an entire email list, for example, or conduct the test on half of your recipients? How long do you want to collect data on the A/B test?

After you answer these questions, you should have all the tools needed to execute an effective A/B test.



Designing A/B Tests

Now that we've covered the principles of A/B tests, let's dive into the details of how to design an A/B test. After all, how do you create multiple versions of the same element, and how do you show different versions to the same audience?

This section helps address these questions. We'll also provide tools and tips to help you design your A/B testing, whether that means tweaking landing pages, calls-to-action, emails, and more.



Building Test Variations

How To Conduct a **Landing Page Test**

With landing page A/B testing you have one URL and two or more versions of the page. When you send traffic to that URL, visitors will be randomly sent to one of your variations. Standard landing page A/B testing tools remember which page the reader landed on and will keep showing that page to the user. For statistical validity split tests, you need to set a cookie on each visitor to ensure the visitor sees the same variation each time they go to the tested page. This is how HubSpot's advanced landing pages tool and Google's Website Optimizer work.

HubSpot's advanced landing pages enable you to create A/B tests and track a number of metrics to evaluate how your experiment is performing. It keeps a record of the number of people who viewed each variation and number of the people who took the intended action. For example, it might inform you that each of your landing page variations was viewed 180 times, with the top-performing one generating 20 clicks and the lowest performing one generating five clicks.

Want a deep dive into the results of your A/B tests once you've split the traffic? Kissmetrics provides an easy interface for you to see the impact of your A/B tests on the rest of your funnel.

While this information is important, it's not enough to make a decision about whether or not your results were significant.

How To Conduct a **Call-to-Action Test**

Call-to-action split testing works pretty much the same way as landing page split testing. You create two variations of your CTA, place them on the same page, and they should be displayed to visitors randomly. The goal here is to determine which call-to-action attracts the most clicks.

However, we mentioned earlier in this ebook that it is important to look for results further down in the sales funnel. This is very easy to do with Kissmetrics's A/B test report. So it will be most useful to know the number of conversions each of your CTA versions drove. This result is influenced by the landing page and how well it is aligned with the call-to-action.

HubSpot's call-to-action module enables you to quickly build A/B tests and drill down to the data that matters the most to your organization. For instance, you might look at the views-to-click rate in an effort to optimize the call-to-action, but if your click-to-submission rate is surprisingly low, then the problem might lie with the landing page. That is why, ultimately, you want to keep an eye on your view-to-submission rate and try to optimize that.

Remember that you should be running only one A/B test at a time, so don't try to optimize both the call-to-action and the landing page simultaneously. Make changes to one variable at a time so you understand which element triggered the results you are seeing.



How To Conduct an Email Test

Most email providers automate the split testing process and enable you to compare different elements about your email. They randomize the list of recipients into two or more groups (you need to ensure the groups are big enough to give you a statistically significant result) and associate each email variation to each of the groups.

HubSpot, for instance, splits your email campaign to help you find out the best subject line and time of day to send an email. The great thing about a tool like this is that it can send the winner to the remainder of your group. This is a fantastic way to optimize your list and deliver the message that attracts the most attention.

HubSpot and most standard email providers enable you to pick a winner based on either open rate or click through rate. However, you also want to see which email is bringing in the most conversions. In other words, identify which variation, combined with the right landing page, delivers the best results. For this type of reporting you need to integrate your email marketing to marketing analytics.



Best Practices for A/B Test Design



Ensure your sample size is large enough

Ideally, you should only conduct A/B tests on a sufficient sample size. This might mean a webpage with high traffic, or a newsletter with a large number of subscribers. If you conduct an A/B test on only a handful of people, for example, these results may not represent your customers on average.



Let the test run sufficiently long enough

An A/B test needs to run for some time to ensure it collects enough data. At the same time, if it runs for too long, seasonality starts to impact the results. Strike a balance between keeping the test's duration long, while ensuring it doesn't go on for too long.



Iterate and optimize future A/B tests

After your first A/B test, take that information and apply to future ones. Why do you think test A performed better than test B? Did the results align with your initial hypothesis, or were they different? Try generating another hypothesis based on these results for *another* A/B test, and then keep going. The more you iterate, the more optimal your marketing efforts will be.



Remember best practices for visuals

Just because you're A/B testing doesn't mean to throw out best practices for visuals. Remember to make them clear and compelling, and ensure they align with your branding. This means including high-quality images, videos, graphics, and so on in all test variations.



Ensure readability across tests

When creating variations on marketing content, remember to ensure readability. This is especially important when updating fonts and colors, as some are more readable than others. For example, light gray text over a white background might be too difficult to read.



Collect user feedback during test

You can also glean insights by asking users directly for feedback, especially if you're A/B testing content such as newsletters and videos. Include a link to a quick survey at the bottom of both email tests, as an example, and see if they comment on the changes.



Avoiding Common Testing Mistakes

One mistake could thwart your A/B test, making results from it inaccurate and unhelpful. Keep these in mind as you build out your A/B tests.



Changing multiple variables at once

It's been stated before, but it's worth repeating: avoid changing multiple variables at once. If you change the color *and* copy of a button, the A/B test is invalid. Ensure there remains *one* variable, otherwise you're not able to determine what caused the change in conversion rates.



Testing on a small audience

The audience needs to be sizable. Otherwise, you run the risk of gleaning insights from an A/B test that's more reflective of the individuals than of your customers.



Projecting your hypothesis onto results

Your hypothesis might not align with the A/B test's results, and that's okay. Try not to project your own bias and perspective onto the data, as this might cause you to view it through an inaccurate perspective.



Paying attention only to quantitative data

An A/B test relies heavily on quantitative data. After all, you're looking at how changes impact conversion rates. That said, qualitative data can offer even more insights into the quantitative data. Remember to ask customers directly when possible. The more data, quantitative and qualitative, the better.



Not optimizing for mobile phones

In all likelihood, your customers will interact with your content on their phones. Undeniably, this will impact your A/B test — especially if you forget to segment your audience between desktop and mobile users.



Remember to not only separate these audiences, but to make your A/B test mobile-friendly. If you change the format of a landing page, make sure it functions on mobile alongside desktop.



Testing on a small audience

If an A/B test shows strong results, you might feel compelled to make the change immediately and permanently. Avoid doing this, as you should run the test a few more times. You never know how different, uncontrollable factors such as seasonality impact data.

Running A/B Tests

Implementing tests correctly

Putting everything together, this section covers how to execute and run your A/B tests. Here's what the process should, ideally, look like:

- ✓ Pick what you want to test (e.g., email, landing page)
- ✓ Define your goals
- ✓ State your hypothesis
- ✓ Select your variable
- ✓ Create two versions, A and B
- ✓ Segment your audience to create a sample size
- ✓ Run the A/B test
- ✓ Extract results and glean insights



Splitting Traffic

Splitting traffic means dividing your audience into two equal, unbiased groups. You can achieve this primarily through randomization.

Typically, most software enables you to randomly split your audience into different groups. If you want to run a newsletter A/B test, for example, you can simply divide your audience randomly into group A and group B, and each group receives their respective variation.

For websites, most website software, including HubSpot's advanced landing pages tool and Google's Website Optimizer, can split traffic that comes to your website into two groups to ensure versions A and B receive an equivalent amount of traffic.

In general, the larger the sample size, the more accurate the results.



Test Duration

The duration of your test depends on a number of factors, including your sample size, seasonality, and traffic volume.

Let's say you want to run an A/B test for your landing pages. You know the page receives, on average, 500 unique visitors per day. If you let the test run for seven days, the page will receive 3,500 visits. In other words, each version will receive 1,750 interactions.

For the most part, this seems like a fairly sizable group to run your A/B test on. At the same time, you may want to run the test again during a different part of the year. Seasonality may impact every result of your A/B test, and you want to reduce these variables as much as possible. The more you run the test at different times of the year, the less impact seasonality has.

Monitoring Tests and Collecting Data

In general, the software behind your newsletters, websites, and so on will automatically collect data for you. That said, you want to figure out which key performance indicators (KPIs) to specifically look at. This might include conversion rates, click-through rates, or bounce rates.

You will have to do some heavy pulling when it comes time to present your results. Create visuals demonstrating the performance of both versions, alongside explanations of how you conducted the A/B test with all of the best practices mentioned previously.

As you monitor the test, keep an eye out for potential errors that may skew the results. Something as simple as a script error, for example, might require you to restart the test.

Maintaining Statistical Significance

When running these tests, you want to maintain statistical significance. That may sound complex, but it's simple. Statistical significance is a way to calculate whether the differences from a test are caused by randomness or an actual variable.

In other words, a *statistically significant* difference means the change (in this case, difference between the performance of versions A and B) cannot be explained by randomness, but by the change in the variable. Statistically significant results means the chances of the observed difference being random are low.

Typically, statistical significance entails a confidence level, which aptly indicates how confident you are in the results. Most tests use a 95% confidence level as default, meaning if you ran the A/B test 100 times, you would expect to observe the difference 95 times.

This kit already includes a calculator you can leverage to more quickly determine statistical significance. For a deeper dive into this, [check out Kissmetrics' full explanation](#).

Analyzing and Interpreting Results

Once you finish the A/B test, you can take a hard look at the selected KPI to extract meaningful insights. Did the results align with your hypothesis, or did they surprise you? Why might that be? Could an external variable, such as seasonality, explain the results? Or is there another hypothesis that might explain the difference in performance?

Interpreting Data Accurately

Of course, you can't immediately take the data at face value. There are other variables to consider, and you want to recognize these and communicate them with other stakeholders.

First, accurately communicate the data you prioritized. Click-through rates? Conversion rates? Bounce rates?

Then, assess the quality of the data. Did the test receive hundreds, or thousands of responses or interactions? Could it have received even more? Did any of the audience demographics skew in a particular direction, suggesting an inadvertent bias?

Additionally, remember to keep in mind the inherent uncertainty in collecting data. You can't control every variable, so acknowledging what could impact the results puts the A/B test in context. Seasonality, a smaller sample size, or even an oversight in experimental design could affect the A/B test.

Finally, keep all of these insights in mind, and start working on repeating the A/B test. By doing this, you limit the effect of those external variables, ensuring you make the right decision.

Deriving Actionable Insights

An A/B test offers many insights you can pull from. Now, attach those insights to actions your organization can take. As you run more A/B tests, for example, you might recommend completely altering your landing page design, or changing the tone of email copy altogether.

As you make these changes, explain the why behind them — even if it's an educated guess. This helps you better characterize your audience, including how they interact with your brand. You might end up finding a difference in performance based on different audience segments. Do younger customers, for example, prefer a different landing page from older ones? What kind of role does the user's device play?

In the end, keep A/B testing. A one-and-done approach might set you up for failure, and continuous A/B testing is a sustainable, long-term way to optimize every part of your marketing efforts.

Using Artificial Intelligence To Drive A/B Tests

Over 80% of industry professionals **plan** to integrate some form of AI into their work. When it comes to A/B testing, that remains the same. AI can greatly improve your A/B tests, whether that means using it as a collaborative tool, a writing tool, or an automation tool.


Generating Hypotheses

You can feed a generative AI tool the context behind your A/B test. Consider the following example, which you can use to get the conversation started.

“

I’m conducting an A/B test with the goal of improving conversion rates on my landing page. It currently has a set of rotating banner images, an introduction to the product, and a blue CTA button. What kind of hypotheses can I generate for the first test?”

The tool will then output potential hypotheses you can pursue. If your organization has an in-house AI tool, with historical data about your business, this can make it align more closely with your long-term strategic goals.



Automating Audience Allocation

Many machine learning tools can also automatically segment your audience for an A/B test. They can take your customer data, analyze demographics, and make segments to personalize your A/B testing. This might include different groups of customers, including their age, buying habits, and so on.

Determining Statistical Significance

Although you can leverage this kit's statistical significance calculator, AI can also automatically do this for you. You simply input the collected data — sample size, conversion rates, and other information — and it can let you know about the confidence levels and p-values to determine statistical significance.

Recommending Changes to Future A/B Tests

As you continue to work with an AI tool, it can iterate on the information it receives to make recommendations for future A/B tests. Let's say it notices a trend — younger users prefer more video content than older ones. It would then recommend further segmentation of your marketing to improve conversion rates.

Creating Visualizations

You can also use AI to create data visualizations when the A/B test finishes. You simply input the collected data and ask for the tool to create visualizations, such as a bar graph comparing the two versions' performances. This saves you significant time, especially if you're working with multiple A/B tests on different parts of the organization.

