

Product is in a pre-release state and might change or have limited support. For more information, see the [product launch stages](#) (</products/#product-launch-stages>).

This page describes how you can use A/B experiments to understand how Recommendations AI is impacting your business.

An [A/B experiment](https://en.wikipedia.org/wiki/A/B_testing) (https://en.wikipedia.org/wiki/A/B_testing) is a randomized experiment with two groups: an experimental group and a control group. The experimental group receives some different treatment (in this case, predictions from Recommendations AI); the control group does not.

When you run an A/B experiment with Recommendations AI, you include the information about which group a user was in when you record user events. Recommendations AI uses that information to refine the model and provide metrics.

Both versions of your application must be the same, except that users in the experimental group see recommendations generated by Recommendations AI and the control group does not. You log user events for both groups.

You set up the experiment using a third-party experiment platform such as [Google Optimize](https://optimize.withgoogle.com) (<https://optimize.withgoogle.com>) or [Optimizely](https://www.optimizely.com) (<https://www.optimizely.com>). The control and experimental groups each get a unique experiment ID from the platform. When you record a user event, you specify which group the user is in by including the experiment ID in the `experimentIds` field of the `eventDetail` section. Providing the experiment ID enables Recommendations AI to compare the metrics for the versions of your application seen by the control and experimental groups.

The goal of an A/B experiment is to accurately determine the impact of updating your site (in this case, employing Recommendations AI). To get an accurate measure of the impact, you must design and implement the experiment correctly, so that other differences do not creep in and impact the experiment results.

To design a meaningful A/B experiment, use the following tips:

- Make sure that the behavior of your site is identical for the experimental group and the control group.

Site behavior includes latency, display format, text format, page layout, image quality, and image size. There should be no discernible differences for any of these attributes between the experience of the control and experiment groups.

- Accept and display recommendations as they are returned from Recommendations AI, and display them in the same order as they are returned.

Filtering out items that are out of stock is acceptable. However, you should avoid filtering or ordering recommendation results based on your business rules.

- Make sure you are including the recommendation token with your user events correctly. [Learn more](#) (/recommendations-ai/docs/recommendation-tokens).
- Make sure that the placement you provide when you request recommendations matches your intention for that recommendation and the location where you display the recommendation results.

The placement affects how models are trained and therefore what products are recommended. [Learn more](#) (/recommendations-ai/docs/placements).

- If you are comparing an existing recommendations solution with Recommendations AI, keep the experience of the control group strictly segregated from the experience of the experimental group.

If the control recommendation solution does not provide a recommendation, do not provide a recommendation from Recommendations AI in the control pages. Doing so will skew your test results.