

Oracle Marketing Recommendations

Increase cross-sells and up-sells by promoting relevant products and content to each customer throughout their journey.

AUTOMATED PERSONALIZATION

Real-time decisions powered by machine learning provide the best experience for each customer.

Leverage the power of machine learning

Oracle Marketing Recommendations makes use of contextual and historical customer data and applies machine learning to this data to identify and serve the right item (content, service, or product) in any given context. The user experience can be tailored automatically at the individual level, starting from the very beginning of the journey from a homepage, a category or listings page, all the way through to product detail pages and the checkout funnel, without the need to define manual business rules.

Test the best recommendation strategy

Marketers can configure advanced strategies by combining machine learning algorithms with a range of filtering options. This enables the recommendation engine's output to tailor the output to each customer even further. Advanced strategies can then be implemented and tested against one another to understand the best approach for the end user. Additionally, marketers can add a layer of targeting, only showing recommendations to certain customers based off of any attribute, whether from CRM, DMP, data layer variables, or onsite behavior captured via JS code.

Increase engagement, up-sell and cross-sell

The primary goal of recommendations is to increase cross-sells and up-sells in a retail environment, or engagement in the case of content. This is achieved by helping customers find what they need or directing their attention to items that they wouldn't have otherwise found by themselves.

Deliver and measure powerful personalization without relying on IT or data science resources

Marketers can inject, configure and launch item recommendations without relying on IT resources. Using our visual Campaign Designer tool, marketers can easily:

- insert a recommendations widget on any part of their site,
- configure the number of slots, information shown per item (title, price, etc.) and the widget's style,
- choose the right machine learning model and apply any suitable filters to the recommendations, target audiences with the most suitable recommendations component (configuration and strategy),
- understand ROI and Recommendations performance through our in-depth campaign reporting.

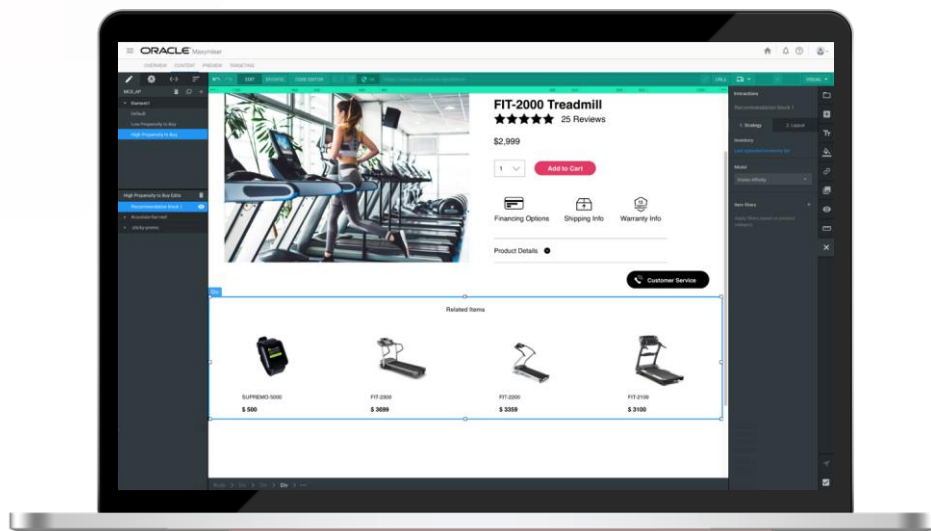
Disclaimer: This document is for informational purposes. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

COMMON USE CASES

Self-serve and easy to use, Oracle Marketing Recommendations empowers marketers to effortlessly inject, configure and deploy real-time recommendations.

Product Recommendations

Marketers in the retail space can leverage product recommendations to tailor the customer experience from the top of the journey to the checkout funnel. For instance, key landing pages (Men's, Women's, etc.) can be personalized based on whether the customer is new or returning, or if they belong to a particular audience with a given propensity to buy from a specific brand. Product recommendations can be served on these landing pages using different strategies for different audiences, e.g., most viewed or best-selling items filtered from the relevant category are suitable for new customers, while last viewed items or items recommended from the visitor affinity model are highly relevant for return customers. For product detail pages, all the standard recommendation strategies are supported, e.g., "People who viewed / bought this item, then went on to view / buy these items."



Content Recommendations

Marketers who want to increase engagement can leverage recommendations to surface relevant and serendipitous content to their users. A key differentiation between use cases is useful here: first, there is the use case of recommending very recent, fast-moving content; and second, the use case of serving content from amongst the most relevant older and more stable items. The first case is often treated more efficiently based on known business rules, e.g., for audiences with known category interests, serve the latest content items from within those categories. The second case is where automated recommendations shine, increasing both the breadth and depth of user engagement at various points across the journey.